



# ***ForeRunner* ASN-9000** **Software Reference Manual**

MANU0272-01 Rev A - November 7, 1997

Software Version ASN\_FT 4.0.0

## **FORE Systems, Inc.**

1000 FORE Drive  
Warrendale, PA 15086-7502  
Phone: 412-742-4444  
FAX: 412-742-7742

<http://www.fore.com>

## Legal Notices

Copyright © 1995-1997 FORE Systems, Inc. All rights reserved. FORE Systems is a registered trademark, and *ForeRunner*, *ForeView*, *ForeThought*, *ForeRunnerLE*, *PowerHub*, and *CellPath* are trademarks of FORE Systems, Inc. All other brands or product names are trademarks of their respective holders.

**U.S. Government Restricted Rights.** If you are licensing the Software on behalf of the U.S. Government (“Government”), the following provisions apply to you. If the Software is supplied to the Department of Defense (“DoD”), it is classified as “Commercial Computer Software” under paragraph 252.227-7014 of the DoD Supplement to the Federal Acquisition Regulations (“DFARS”) (or any successor regulations) and the Government is acquiring only the license rights granted herein (the license rights customarily provided to non-Government users). If the Software is supplied to any unit or agency of the Government other than DoD, it is classified as “Restricted Computer Software” and the Government’s rights in the Software are defined in paragraph 52.227-19 of the Federal Acquisition Regulations (“FAR”) (or any successor regulations) or, in the cases of NASA, in paragraph 18.52.227-86 of the NASA Supplement to the FAR (or any successor regulations).

Printed in the USA.

No part of this work covered by copyright may be reproduced in any form. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

This publication is provided by FORE Systems, Inc. “as-is” without warranty of any kind, either express or implied, including, but not limited to, the implied warranties or conditions of merchantability or fitness for a particular purpose. FORE Systems, Inc. shall not be liable for any errors or omissions which may occur in this publication, nor for incidental or consequential damages of any kind resulting from the furnishing, performance, or use of this publication.

Information published here is current or planned as of the date of publication of this document. Because we are improving and adding features to our products continuously, the information in this document is subject to change without notice.

RESTRICTED RIGHTS LEGEND. Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 (October 1988) and FAR 52.227-19 (June 1987).

The VxWorks software used in the Mini Loader is licensed from Wind River Systems, Inc., Copyright ©1984-1996.

## FCC CLASS A NOTICE

**WARNING:** Changes or modifications to this unit not expressly approved by the party responsible for compliance could void this user’s authority to operate this equipment.

NOTE: The *ForeRunner* ASN-9000 has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15, FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

## DOC CLASS A NOTICE

This digital apparatus does not exceed Class A limits for radio noise emission for a digital device as set out in the Radio Interference Regulations of the Canadian Department of Communications.

Le present appareil numerique n’emet pas de bruits radioelectriques depassant les limites applicables aux appareils numeriques de la class A prescrites dans le reglement sur le brouillage radioelectrique edicte par le ministere des Communications du Canada.

## **VCCI CLASS 1 NOTICE**

この装置は、第一種情報処理装置（商工業地域において使用されるべき情報処理装置）で商工業地域での電波障害防止を目的とした情報処理装置等電波障害自主規制協議会(VCCI)基準に適合しております。

従って、住宅地域またはその隣接した地域で使用する、ラジオ、テレビジョン受信機等に受信障害を与えることがあります。

取扱説明書に従って正しい取り扱いをして下さい。

This equipment is in the Class 1 category (Information Technology Equipment to be used in commercial and/or industrial areas) and conforms to the standards set by the Voluntary Control Council For Interference by Information Technology Equipment aimed at preventing radio interference in commercial and/or industrial areas. Consequently, when used in a residential area or in an adjacent area thereto, radio interference may be caused to radios and TV receivers, etc. Read the instructions for correct handling.

## **CE NOTICE**

Marking by the symbol **CE** indicates compliance of this system to the EMC (Electromagnetic Compatibility) directive of the European Community and compliance to the Low Voltage (Safety) Directive. Such marking is indicative that this system meets or exceeds the following technical standards:

- EN 55022 - "Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment."
- EN 50082-1 - "Electromagnetic compatibility - Generic immunity standard Part 1: Residential, commercial, and light industry."

## **SAFETY CERTIFICATIONS**

ETL certified to meet Information Technology Equipment safety standards UL 1950 3rd Edition, CSA22.2, No. 950-95, EN 60950 (1992) and IEC 950, 2nd Edition.

## CANADIAN IC CS-03 COMPLIANCE STATEMENT

**NOTICE:** The Industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational and safety requirements. The Industry Canada label does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. In some cases, the company's inside wiring associated with a single line individual service may be extended by means of a certified connector assembly (telephone extension cord). The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs to certified equipment should be made by an authorized Canadian maintenance facility designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important in rural areas.

**Caution:** Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

## TRADEMARKS

FORE Systems is a registered trademark, and *ForeView* and *ForeRunner* are trademarks of FORE Systems, Inc. All other brands or product names are trademarks or registered trademarks of their respective holders.

# Table of Contents

**Table of Contents**

**List of Figures**

**List of Tables**

**Preface**

Chapter Summaries.....	i
Related Publications .....	iii
Technical Support .....	iii
Typographical Styles .....	iv
Important Information Indicators .....	v
Laser Notice .....	vi
Safety Precautions.....	vii
Modifications to Equipment .....	vii
Placement of a FORE Systems Product .....	vii
Power Cord Connection .....	viii

**CHAPTER 1   Features Overview**

1.1	Intelligent Packet Switching .....	1 - 1
1.2	Software Features .....	1 - 1
1.2.1	System Management Features .....	1 - 2
1.2.1.1	Multiprocessor Optimization .....	1 - 2
1.2.1.2	Multiple Boot Sources .....	1 - 2
1.2.1.3	UNIX-Like Command-Line Interface .....	1 - 3
1.2.1.4	Local File-Management System .....	1 - 3
1.2.1.5	Concurrent Command-Line Sessions .....	1 - 3
1.2.1.6	System Configuration Files .....	1 - 3
1.2.1.7	Session Parameter Files .....	1 - 3
1.2.1.8	Automatic Segment-State Detection .....	1 - 4
1.2.1.9	Segment Statistics.....	1 - 4
1.2.2	Virtual Local Area Networks (VLANs) .....	1 - 4
1.2.3	Bridging and Routing Features.....	1 - 4
1.2.3.1	Bridge Table and Bridge Cache.....	1 - 4
1.2.3.1.1	802.1d.....	1 - 5
1.2.3.1.2	Spanning-Tree .....	1 - 5
1.2.3.1.3	IPX Translation Bridging .....	1 - 5
1.2.3.2	IP Routing .....	1 - 5

## Table of Contents

	1.2.3.2.1	Routing Information Protocol (RIP) . . . . .	1 - 6
	1.2.3.2.2	Open Shortest Path First (OSPF) . . . . .	1 - 6
	1.2.3.3	AppleTalk Routing. . . . .	1 - 6
	1.2.3.4	IPX Routing . . . . .	1 - 6
	1.2.3.5	DECnet Routing . . . . .	1 - 6
	1.2.3.6	Route Protocol Statistics . . . . .	1 - 7
	1.2.3.7	Security Filters . . . . .	1 - 7
1.3		Network Management Features . . . . .	1 - 7
	1.3.1	Network Management System (NMS). . . . .	1 - 7
	1.3.2	Management Information Base (MIB) Agents . . . . .	1 - 8
	1.3.3	ForeView . . . . .	1 - 8

## CHAPTER 2 ASN-9000 Files

2.1	Types . . . . .	2 - 1
	2.1.1 Software Files . . . . .	2 - 2
2.2	Other Files Supplied with the ASN-9000 . . . . .	2 - 2
	2.2.1 Files the ASN-9000 Creates . . . . .	2 - 3

## CHAPTER 3 Using the Command Line Interface

3.1	Using the New User Interface . . . . .	3 - 1
3.1.1	Issuing Commands . . . . .	3 - 1
3.1.2	Entering and Editing Command Text. . . . .	3 - 3
3.1.3	Displaying a Command's Syntax. . . . .	3 - 3
3.1.4	Command Syntax . . . . .	3 - 4
3.1.4.1	Verbs . . . . .	3 - 4
3.1.4.1.1	set and unset . . . . .	3 - 4
3.1.4.1.2	define and undefine . . . . .	3 - 4
3.1.4.1.3	attach and detach . . . . .	3 - 4
3.1.4.1.4	add and delete. . . . .	3 - 4
3.1.4.1.5	enable and disable. . . . .	3 - 5
3.1.4.1.6	show and clear. . . . .	3 - 5
3.1.4.2	Objects . . . . .	3 - 5
3.1.4.2.1	config. . . . .	3 - 5
3.1.4.2.2	status. . . . .	3 - 5
3.1.4.2.3	stats. . . . .	3 - 5
3.1.4.2.4	interface. . . . .	3 - 5
3.1.4.2.5	route . . . . .	3 - 5
3.1.4.2.6	bt . . . . .	3 - 5
3.1.4.2.7	cache. . . . .	3 - 6
3.1.4.3	Parameters . . . . .	3 - 6
3.1.4.3.1	Keyword Parameters . . . . .	3 - 6
3.1.4.3.2	Positional Parameters . . . . .	3 - 6

**CHAPTER 4 Global Commands**

4.1	Rebooting the Switch . . . . .	4 - 1
4.1.1	Reset Switch . . . . .	4 - 1
4.1.2	Using the Reboot Command . . . . .	4 - 2
4.1.3	Booting the ASN-9000 . . . . .	4 - 3
4.2	Logging In and Out . . . . .	4 - 3
4.2.1	Logging In . . . . .	4 - 4
4.2.1.1	TTY . . . . .	4 - 4
4.2.1.1.1	Baud Rates . . . . .	4 - 4
4.2.1.2	TELNET . . . . .	4 - 5
4.2.2	Logging Out . . . . .	4 - 5
4.3	Changing Management Capability . . . . .	4 - 6

**CHAPTER 5 More Global Commands**

5.1	Displaying and Using Command Histories . . . . .	5 - 1
5.1.1	Resetting the Return Code . . . . .	5 - 2
5.2	Managing Files . . . . .	5 - 2
5.2.1	Displaying a Directory . . . . .	5 - 2
5.2.2	Deleting a File . . . . .	5 - 3
5.3	Using Command Aliases . . . . .	5 - 4
5.3.1	Defining an Alias . . . . .	5 - 4
5.3.2	Displaying an Alias . . . . .	5 - 5
5.3.3	Saving and Loading an Alias . . . . .	5 - 5
5.3.4	Deleting an Alias . . . . .	5 - 6
5.4	Using Timed Commands . . . . .	5 - 6
5.4.1	Defining a Timed Command . . . . .	5 - 7
5.4.2	Starting a Timed Command . . . . .	5 - 8
5.4.3	Stopping a Timed Command . . . . .	5 - 8
5.4.4	Deleting a Timed Command . . . . .	5 - 9
5.5	Using Environment Files . . . . .	5 - 9
5.5.1	Saving an Environment File . . . . .	5 - 11
5.5.2	Reading (Loading) an Environment File . . . . .	5 - 13
5.5.3	Editing an Environment File . . . . .	5 - 13

**CHAPTER 6 System Subsystem Commands**

6.1	Accessing the System Subsystem . . . . .	6 - 1
6.2	System Commands . . . . .	6 - 2
6.2.1	Changing the TTY Port Baud Rate . . . . .	6 - 2
6.2.2	Removing and Replacing Interface Modules . . . . .	6 - 3
6.2.3	Displaying the ASN-9000 . . . . .	6 - 5
6.2.4	Setting and Displaying the System Time and Date . . . . .	6 - 6

## Table of Contents

6.2.5	Setting the Data Carrier Detect Parameter . . . . .	6 - 7
6.2.6	Displaying the ASN-9000 MAC Address . . . . .	6 - 8
6.2.7	Displaying ID and Power Information . . . . .	6 - 8
6.2.8	Changing the Password. . . . .	6 - 9
6.2.9	Reading a Configuration File. . . . .	6 - 10
6.2.9.1	Loading a Configuration from Flash Memory . . . . .	6 - 11
6.2.9.2	Loading a Configuration From a TFTP Server . . . . .	6 - 11
6.2.10	Rebooting the ASN-9000 . . . . .	6 - 12
6.2.11	Saving the Configuration . . . . .	6 - 12
6.2.11.1	Saving Configuration Files to Flash Memory . . . . .	6 - 12
6.2.11.2	Saving the Configuration to a TFTP Server . . . . .	6 - 12
6.2.12	Setting and Displaying the System Location . . . . .	6 - 14
6.2.13	Setting and Displaying the System Name . . . . .	6 - 14
6.2.14	Displaying the Temperature of the ASN-9000 . . . . .	6 - 15
6.2.15	Enabling/Disabling TTY2. . . . .	6 - 16
6.2.16	Displaying the System Uptime. . . . .	6 - 16
6.2.17	Displaying the ASN-9000 Software Version . . . . .	6 - 17
6.2.18	Displaying Boot Information on the ASN-9000 . . . . .	6 - 18
6.3	System Configuration Commands. . . . .	6 - 18
6.3.1	Rebooting Without Loading the Default Configuration File . . . . .	6 - 18
6.3.2	Editing a Configuration File . . . . .	6 - 19
6.3.3	Capturing Configuration Information . . . . .	6 - 19
<b>CHAPTER 7 NVRAM Subsystem Commands</b>		
7.1	NVRAM Configuration Commands . . . . .	7 - 1
7.1.1	Boot Order. . . . .	7 - 2
7.1.2	My IP Address. . . . .	7 - 2
7.1.3	My Subnet Mask . . . . .	7 - 3
7.1.4	File Server IP Address . . . . .	7 - 3
7.1.5	Gateway IP Address . . . . .	7 - 3
7.1.6	Crash Reboot . . . . .	7 - 4
7.1.7	Slot Segments. . . . .	7 - 4
7.2	RIPv2 Authentication . . . . .	7 - 5
<b>CHAPTER 8 Media Subsystem Commands</b>		
8.1	Displaying Bridge-Related Configuration. . . . .	8 - 2
8.2	Inter-Segment Statistics . . . . .	8 - 3
8.3	Ethernet LED Modes . . . . .	8 - 4
8.4	Operating-Mode. . . . .	8 - 4
8.5	UTP Port Receiver Status . . . . .	8 - 4
8.6	Displaying Port-Level Statistics . . . . .	8 - 4
8.7	Configuring Packet Forwarding on Segments . . . . .	8 - 5

8.8	Configuring Segment Names . . . . .	8 - 5
8.9	Automatic Segment-State Detection. . . . .	8 - 6
8.9.1	Software Behavior When Disabled. . . . .	8 - 7
8.9.2	Default Setting . . . . .	8 - 7
8.10	Setting Segment-State Threshold . . . . .	8 - 7
8.11	Status . . . . .	8 - 7
8.12	Media-Level Statistics. . . . .	8 - 8

## **CHAPTER 9 Host Subsystem Commands**

9.1	Accessing the Host Subsystem . . . . .	9 - 1
9.2	Displaying the TCP Configuration. . . . .	9 - 2
9.3	Displaying the TCP Table . . . . .	9 - 3
9.4	Displaying Statistics . . . . .	9 - 4
9.4.1	Clearing Statistics. . . . .	9 - 5
9.5	Setting TCP Session Parameters. . . . .	9 - 5
9.5.1	The Connection Idle Time and Keep Alive Interval . . . . .	9 - 5
9.5.1.1	Setting the Connection Idle Time . . . . .	9 - 6
9.5.1.2	Setting the Keep-alive Interval. . . . .	9 - 6
9.5.2	Closing a TCP Connection . . . . .	9 - 6
9.5.2.1	Current TCP Connection . . . . .	9 - 6
9.5.2.2	Another TCP Connection . . . . .	9 - 6
9.6	Displaying the UDP Table . . . . .	9 - 7

## **CHAPTER 10 TFTP Subsystem Commands**

10.1	Accessing the TFTP Subsystem . . . . .	10 - 1
10.2	Considerations . . . . .	10 - 2
10.2.1	TFTP Commands and UNIX Read/Write Permissions. . . . .	10 - 2
10.2.2	PathNames. . . . .	10 - 3
10.2.3	File Naming Conventions . . . . .	10 - 4
10.2.4	Remote File Names . . . . .	10 - 4
10.3	Setting, Displaying, or Unsetting the Default Server. . . . .	10 - 5
10.3.1	Setting the Default TFTP Server . . . . .	10 - 5
10.3.2	Displaying the Default TFTP Server . . . . .	10 - 5
10.3.3	Removing the Default TFTP Server Setting . . . . .	10 - 5
10.4	Downloading or Displaying a File . . . . .	10 - 6
10.5	Uploading a File . . . . .	10 - 7
10.6	Loading a Configuration File. . . . .	10 - 8
10.7	Saving a Configuration File. . . . .	10 - 9

## CHAPTER 11 Bridge Subsystem Commands

11.1	Accessing the Bridge Subsystem .....	11 - 1
11.2	Showing the Bridging Configuration .....	11 - 1
11.3	Using the Bridge Table. ....	11 - 4
11.3.1	Displaying the Bridge Table .....	11 - 4
11.3.2	Clearing the Bridge Table .....	11 - 6
11.3.3	Adding an Entry to the Bridge Table .....	11 - 6
11.3.4	Enabling and Disabling Bridge Learning .....	11 - 8
11.3.5	Changing the Aging Interval .....	11 - 8
11.3.6	Deleting an Entry from the Bridge Table .....	11 - 9
11.4	Defining and Adding Bridge Groups .....	11 - 9
11.4.1	Displaying the Bridge Groups .....	11 - 10
11.4.2	Deleting a Bridge Group .....	11 - 11
11.5	Displaying the Bridging Status of a Segment .....	11 - 11
11.6	Configuring Spanning-Tree Parameters .....	11 - 13
11.6.1	Enabling or Disabling Spanning Tree .....	11 - 14
11.6.2	Changing Spanning-Tree Parameters .....	11 - 14
11.6.2.1	Setting the Bridge Priority. ....	11 - 14
11.6.2.2	Setting a Segment's Priority .....	11 - 15
11.6.2.3	Setting the Path Cost .....	11 - 15
11.6.2.4	Setting the Maximum Age .....	11 - 16
11.6.2.5	Setting the Hello Time .....	11 - 16
11.6.2.6	Setting the Forward Delay .....	11 - 16
11.6.2.7	Setting the Fast-Hello Time .....	11 - 17
11.6.2.8	Setting the High- and Low-Utilization Percentage .....	11 - 17
11.7	Displaying and Clearing Bridge Statistics .....	11 - 18
11.7.1	Clearing Statistics .....	11 - 18
11.8	Displaying and Clearing the Bridge Cache .....	11 - 18
11.8.1	Displaying the Bridge Cache .....	11 - 19
11.8.2	Clearing the Bridge Cache .....	11 - 19

## CHAPTER 12 SNMP Subsystem Commands

12.1	Accessing the SNMP Subsystem .....	12 - 1
12.2	Displaying the SNMP Configuration .....	12 - 2
12.3	Displaying Statistics. ....	12 - 2
12.3.1	Clearing Statistics .....	12 - 3
12.4	Adding an SNMP Community .....	12 - 3
12.4.1	Supported SNMP Traps. ....	12 - 4
12.4.2	Deleting an SNMP Community .....	12 - 4
12.4.3	Adding an SNMP Manager .....	12 - 4

12.4.4	Deleting an SNMP Manager . . . . .	12 - 5
12.4.5	Preparing Files for SunNet Manager . . . . .	12 - 5

## CHAPTER 13 IP Subsystem Commands

13.1	Accessing the IP Subsystem . . . . .	13 - 1
13.2	Displaying the IP Configuration . . . . .	13 - 2
13.3	Configuring and Showing IP Interfaces . . . . .	13 - 3
13.3.1	Considerations . . . . .	13 - 4
13.3.2	Restrictions. . . . .	13 - 5
13.3.3	How the Software Handles IP Packets . . . . .	13 - 5
13.3.4	Showing the IP Interface Table . . . . .	13 - 6
13.3.5	Adding an IP Interface . . . . .	13 - 6
13.3.6	Deleting an IP Interface . . . . .	13 - 9
13.3.7	Configuring VLANs. . . . .	13 - 9
13.3.7.1	Changing the VLAN Configuration . . . . .	13 - 10
13.3.7.2	Deleting a Configured VLAN . . . . .	13 - 11
13.3.8	Allocating Memory for Additional IP Routes . . . . .	13 - 12
13.3.9	Enabling IP Routing . . . . .	13 - 12
13.4	Showing, Adding, and Deleting IP Routes . . . . .	13 - 12
13.4.1	Showing the IP Route Table . . . . .	13 - 13
13.4.2	Adding an IP Route . . . . .	13 - 15
13.4.3	Enabling and Disabling Load Balancing . . . . .	13 - 16
13.4.4	Enabling Loopback Detection. . . . .	13 - 16
13.4.4.1	Setting the Loopback Detection Time . . . . .	13 - 16
13.4.4.2	Displaying the IP Loop Detection Table . . . . .	13 - 17
13.4.5	Enabling or Disabling an IP Route . . . . .	13 - 17
13.4.6	Deleting an IP Route . . . . .	13 - 18
13.5	IP Router Discovery . . . . .	13 - 18
13.5.1	Setting the Advertisement Address . . . . .	13 - 19
13.5.2	Setting the Advertisement Preference . . . . .	13 - 20
13.5.3	Setting the Advertisement Interval . . . . .	13 - 20
13.5.4	Displaying the Advertisement Interval . . . . .	13 - 21
13.6	Showing and Configuring the ARP Table . . . . .	13 - 21
13.6.1	Enabling and Disabling ARP . . . . .	13 - 21
13.6.2	The ARP Cache . . . . .	13 - 22
13.6.3	Showing the ARP Table . . . . .	13 - 22
13.6.4	Clearing the ARP Table . . . . .	13 - 23
13.6.5	Showing and Changing the ARP Aging Interval . . . . .	13 - 23
13.6.6	Adding a Static Entry to the ARP Table . . . . .	13 - 24
13.6.7	Deleting a Static Entry from the ARP Table . . . . .	13 - 25
13.7	Pinging Other IP Devices . . . . .	13 - 25

## Table of Contents

13.8	IP Helper .....	13 - 26
13.8.1	How IP Helper Works .....	13 - 27
13.8.2	Using IP Helper .....	13 - 28
13.8.2.1	Adding an IP Helper Address .....	13 - 28
13.8.2.2	Deleting an IP Helper Address .....	13 - 29
13.8.2.3	Displaying Statistics and the UDP Table .....	13 - 29
13.8.2.4	Deleting Default UDP Entries .....	13 - 30
13.8.2.5	Clearing Statistics .....	13 - 30
13.8.2.6	Deleting an IP Helper Address .....	13 - 31
13.8.2.7	Adding an IP Helper Gateway IP Address .....	13 - 31
13.8.2.8	Deleting an IP Helper Gateway Address .....	13 - 32
13.8.2.9	Displaying IP Helper Gateway Addresses .....	13 - 32
13.8.3	Setting the Time-To-Live Parameter .....	13 - 33
13.8.4	Enabling and Disabling ICMP Redirect Messages .....	13 - 33
13.8.5	Enabling or Disabling Source-Route Filtering .....	13 - 33
13.8.6	Enabling or Disabling Network-Broadcast Forwarding .....	13 - 34
13.8.6.1	Disabling Bridging of Net Broadcasts .....	13 - 36
13.8.6.2	Disabling Routing of Net Broadcasts .....	13 - 36
13.8.7	Enabling Proxy ARP .....	13 - 36
13.8.7.1	Displaying the Proxy ARP Table .....	13 - 37
13.9	Showing and Clearing Statistics .....	13 - 38
13.9.1	Clearing Statistics .....	13 - 39
13.10	Showing or Clearing the IP Route Cache .....	13 - 39
13.10.1	Displaying the Route Cache .....	13 - 39
13.10.2	Flushing the Route Cache .....	13 - 40

## CHAPTER 14 IP Multicast Subsystem Commands

14.1	Accessing the IP Multicast Subsystem .....	14 - 1
14.1.1	Allocating Memory .....	14 - 2
14.1.2	Enabling Pruning .....	14 - 2
14.2	Showing the IP Multicast Configuration .....	14 - 3
14.2.1	IP Considerations .....	14 - 3
14.2.2	Displaying IP Multicast Groups .....	14 - 4
14.2.3	Displaying IP Multicast Neighbors .....	14 - 4
14.3	Configuring and Showing IP Multicast Interfaces .....	14 - 5
14.3.1	Displaying the Interface Table .....	14 - 6
14.3.2	Deleting a Physical Interface .....	14 - 7
14.3.3	Deleting a Tunnel .....	14 - 7
14.4	Configuring and Showing Tunnels .....	14 - 8
14.4.1	Adding a Tunnel .....	14 - 8
14.4.2	Deleting a Tunnel .....	14 - 9

14.5	Enabling IP Multicast Routing . . . . .	14 - 9
14.5.1	Enabling Multicast Traffic on a Segment . . . . .	14 - 10
14.6	Configuring and Showing IP Multicast Routes . . . . .	14 - 10
14.6.1	Clearing the Route Table . . . . .	14 - 12
14.7	Using the IP Route Cache . . . . .	14 - 12
14.7.1	Displaying and Clearing the Route Cache . . . . .	14 - 12
14.8	Displaying Statistics . . . . .	14 - 12
14.8.1	Clearing Statistics. . . . .	14 - 15
14.9	Enabling Multicast-Aware Bridging . . . . .	14 - 15

## **CHAPTER 15 IP/RIP Subsystem Commands**

15.1	Accessing the RIP Subsystem . . . . .	15 - 1
15.2	Displaying the RIP Configuration . . . . .	15 - 1
15.2.1	Configuring RIP Parameters. . . . .	15 - 2
15.2.2	Enabling Acceptance of Default Routes . . . . .	15 - 4
15.2.3	Enabling Authentication of RIP Updates . . . . .	15 - 4
15.2.4	Setting the Authorization String on a VLAN . . . . .	15 - 5
15.2.5	Enabling Report of Learned Routes. . . . .	15 - 5
15.2.6	Setting the Receive and Transmit Type on a VLAN . . . . .	15 - 6
15.3	Displaying and Clearing RIP Statistics . . . . .	15 - 7
15.4	Bridging RIP Updates Across VLANs. . . . .	15 - 7

## **CHAPTER 16 IP/OSPF Subsystem Commands**

16.1	Accessing the IP/OSPF Subsystem . . . . .	16 - 1
16.2	Configuring an ASN-9000 Switch as an OSPF Router. . . . .	16 - 1
16.2.1	Allocating Memory . . . . .	16 - 2
16.2.2	Assigning the OSPF Router ID. . . . .	16 - 2
16.2.3	Displaying the Router-ID . . . . .	16 - 3
16.2.4	Enabling OSPF. . . . .	16 - 4
16.2.5	Enabling the ASN-9000 as a System Border Router . . . . .	16 - 4
16.2.6	Setting the Automatic Virtual-Link Feature. . . . .	16 - 5
16.2.6.1	Displaying the Virtual-Link Table . . . . .	16 - 5
16.2.7	Adding an OSPF Interface to an Area . . . . .	16 - 6
16.2.8	Using the NSET Command . . . . .	16 - 6
16.2.9	Adding an OSPF Area . . . . .	16 - 9
16.2.9.1	Deleting an OSPF Area. . . . .	16 - 11
16.2.10	Displaying an OSPF Area. . . . .	16 - 11
16.2.11	Adding Network Ranges. . . . .	16 - 13
16.2.12	Deleting Network Ranges. . . . .	16 - 14
16.2.13	Displaying Network Ranges . . . . .	16 - 14
16.2.14	Displaying OSPF Neighbors. . . . .	16 - 15

## Table of Contents

16.2.15	Displaying OSPF Link-State Advertisements . . . . .	16 - 17
16.2.16	Enabling the Return-Code Prompt . . . . .	16 - 20
16.2.17	Adding a Virtual-Link . . . . .	16 - 20
16.2.18	Deleting a Virtual-Link . . . . .	16 - 21
16.2.19	Displaying Virtual-Links . . . . .	16 - 22
16.2.20	Timed Commands . . . . .	16 - 26
16.2.21	Statistics Command . . . . .	16 - 26
16.2.22	Displaying OSPF Statistics . . . . .	16 - 26
16.2.23	Clearing OSPF Statistics . . . . .	16 - 27
 <b>CHAPTER 17 AppleTalk Subsystem Commands</b>		
17.1	Accessing the AppleTalk Subsystem . . . . .	17 - 2
17.2	Getting Started . . . . .	17 - 2
17.2.1	Enabling the AppleTalk Subsystem . . . . .	17 - 2
17.2.1.1	Allocating Memory . . . . .	17 - 2
17.2.1.2	Enabling AppleTalk Routing . . . . .	17 - 3
17.2.1.3	Displaying the Current Configuration . . . . .	17 - 3
17.2.1.4	Saving Your AppleTalk Configuration . . . . .	17 - 4
17.3	Configuring ASN-9000 Segments for AppleTalk . . . . .	17 - 4
17.3.1	The Zone Commands . . . . .	17 - 4
17.3.1.1	Adding a Zone Name . . . . .	17 - 4
17.3.1.2	Displaying the Zone Information . . . . .	17 - 6
17.3.1.2.1	Configured Zones . . . . .	17 - 6
17.3.1.3	Deleting a Configured Zone . . . . .	17 - 7
17.4	Configuring AppleTalk Interfaces . . . . .	17 - 8
17.4.1	Adding a Interface (Network Address) . . . . .	17 - 8
17.4.2	Displaying Network Address Information . . . . .	17 - 10
17.4.3	Deleting a Network Address . . . . .	17 - 13
17.5	Using the AARP Table . . . . .	17 - 15
17.5.1	Displaying AARP Entries . . . . .	17 - 16
17.5.2	Setting the AARP Aging Time . . . . .	17 - 17
17.5.3	Clearing the AARP Table . . . . .	17 - 17
17.6	Displaying Route Information . . . . .	17 - 17
17.7	Using the Route Cache . . . . .	17 - 20
17.7.1	Displaying the Route Cache . . . . .	17 - 20
17.7.2	Flushing the Route Cache . . . . .	17 - 20
17.8	Displaying NBP Information . . . . .	17 - 21
17.9	Displaying Statistics . . . . .	17 - 21
17.10	Clearing AppleTalk Statistics . . . . .	17 - 23
17.11	Testing a Network Address . . . . .	17 - 23

**CHAPTER 18 IPX Subsystem Commands**

18.1	Accessing the IPX Subsystem . . . . .	18 - 1
18.2	Allocating Memory for IPX Routing. . . . .	18 - 1
18.3	Showing the IPX Configuration. . . . .	18 - 2
18.4	Adding and Deleting IPX Interfaces . . . . .	18 - 3
18.4.1	Deleting IPX Interfaces. . . . .	18 - 5
18.5	Displaying IPX Interfaces . . . . .	18 - 5
18.6	Enabling IPX Routing . . . . .	18 - 6
18.6.1	Adding and Deleting IPX Routes . . . . .	18 - 6
18.6.2	Deleting IPX Routes . . . . .	18 - 8
18.6.3	Displaying IPX Routes . . . . .	18 - 8
18.7	Displaying and Clearing the IPX Route Cache. . . . .	18 - 10
18.7.1	Displaying the Route Cache . . . . .	18 - 10
18.7.2	Clearing the Route Cache . . . . .	18 - 12
18.8	Configuring IPX RIP and SAP Parameters. . . . .	18 - 13
18.8.1	Setting the Control Type . . . . .	18 - 13
18.8.1.1	Displaying the RIP and SAP Control Type. . . . .	18 - 14
18.8.1.2	Adjusting the Interval and Aging Timers . . . . .	18 - 14
18.8.2	Setting the RIP Parameters . . . . .	18 - 15
18.8.3	Setting the SAP Parameters. . . . .	18 - 16
18.8.4	Displaying the Configuration. . . . .	18 - 17
18.8.5	Setting the Parameters. . . . .	18 - 17
18.8.6	Equal RIP Route. . . . .	18 - 18
18.9	Using the Server Table . . . . .	18 - 18
18.9.1	Displaying the Server Table . . . . .	18 - 19
18.9.2	Adding a Static Server . . . . .	18 - 21
18.9.3	Deleting a Static Server . . . . .	18 - 22
18.10	Using IPX Helper . . . . .	18 - 22
18.10.1	Adding an IPX Helper Address. . . . .	18 - 23
18.10.2	Displaying an IPX Helper . . . . .	18 - 23
18.10.3	Deleting an IPX Helper Address. . . . .	18 - 23
18.11	Showing and Clearing Statistics. . . . .	18 - 24
18.12	Customizing the IPX Configuration. . . . .	18 - 25
18.12.1	Type-20 Forwarding for Segments . . . . .	18 - 25
18.12.2	Enabling Large Packets . . . . .	18 - 25

**CHAPTER 19 Configuring IPX Translation Bridging**

19.1	Encapsulation Types. . . . .	19 - 1
19.2	Configuration Requirements. . . . .	19 - 2
19.2.1	Enabling IPX Translation Bridging . . . . .	19 - 2

## Table of Contents

19.2.2	Adding IPX Translation-Bridging Interfaces . . . . .	19 - 2
19.2.3	Displaying IPX Translation-Bridging Interfaces . . . . .	19 - 3
19.2.4	Deleting IPX Translation-Bridging Interfaces . . . . .	19 - 4
<b>CHAPTER 20 DECnet Subsystem Commands</b>		
20.1	Accessing the DECnet Subsystem . . . . .	20 - 1
20.1.1	Allocating Memory . . . . .	20 - 2
20.1.2	Node Configuration . . . . .	20 - 2
20.1.3	DECnet Network Topology Restrictions . . . . .	20 - 4
20.1.4	Configuring the ASN-9000 as a DECnet Node . . . . .	20 - 4
20.1.4.1	Additional Node Commands . . . . .	20 - 6
20.2	Segment Configuration . . . . .	20 - 9
20.2.1	Configuration . . . . .	20 - 9
20.3	Display Commands . . . . .	20 - 11
20.3.1	Verification of Routing . . . . .	20 - 11
20.3.2	Setting and Displaying Block-Size . . . . .	20 - 13
20.3.3	Displaying Adjacent Routers . . . . .	20 - 13
20.3.4	Displaying Adjacent Endnodes . . . . .	20 - 14
20.3.5	Displaying the Route Table . . . . .	20 - 14
20.3.6	Displaying Statistics . . . . .	20 - 17
20.3.6.1	Displaying the Route Cache . . . . .	20 - 18
<b>APPENDIX A Configuration Defaults</b>		
A.1	ASN-9000 Software Subsystems and Defaults . . . . .	A - 1
<b>APPENDIX B Well-Known Ports</b>		
<b>Index</b>		

# List of Figures

Figure 3.1	User Interface Command Line . . . . .	3 - 1
Figure 5.1	Details List of the Flash Memory Module . . . . .	5 - 3
Figure 15.1	Example of RIP Bridging . . . . .	15 - 8
Figure 17.1	<b>interface show</b> command details. . . . .	17 - 11
Figure 20.1	Illegal Double Links. . . . .	20 - 4
Figure 20.2	Router Verification . . . . .	20 - 12

## *List of Figures*

# List of Tables

Table 5.1	Timed Commands . . . . .	5 - 7
Table 8.1	Segment Level Statistic Parameters . . . . .	8 - 9
Table 11.1	<b>config</b> Command Arguments . . . . .	11 - 2
Table 12.1	SunNet Manager utilities.. . . . .	12 - 6
Table 16.1	LSA Types. . . . .	16 - 18
Table 18.1	Server-Type Mnemonics . . . . .	18 - 20
Table 19.1	Encapsulation types . . . . .	19 - 1
Table A.1	Bridge Subsystem . . . . .	A - 1
Table A.2	Global Commands . . . . .	A - 2
Table A.3	Host Subsystem . . . . .	A - 2
Table A.4	TFTP Subsystem . . . . .	A - 3
Table A.5	IP Subsystem . . . . .	A - 3
Table A.6	IP Multicast Subsystem . . . . .	A - 4
Table A.7	IP/RIP Subsystem. . . . .	A - 5
Table A.8	IP/OSPF Subsystem. . . . .	A - 6
Table A.9	ATALK Subsystem . . . . .	A - 6
Table A.10	IPX Subsystem . . . . .	A - 7
Table A.11	DECnet Subsystem . . . . .	A - 7
Table B.1	Well Known Names and Ports . . . . .	B - 2

## *List of Tables*

# Preface

This manual describes the *ForeRunner* ASN-9000 user interface and the software commands you use to configure and manage the *ForeRunner* ASN-9000 switch for bridging and routing. To learn how to create and apply filters to control the traffic received and forwarded by the *ForeRunner* ASN-9000 switch, see the *PowerHub Filters Manual*.

## Chapter Summaries

---

**Chapter 1 - Features Overview** - Describes the ASN-9000 software features.

**Chapter 2 - ASN-9000 Files** - Describes the files that are shipped with the ASN-9000 and the ASN-9000 can create.

**Chapter 3 - Using the Command line Interface** - Describes the commands in the system subsystem.

**Chapter 4 - Global Commands** - Describes the Global Commands which are available from all subsystems within the ASN-9000.

**Chapter 5 - More Global Commands** - Describes the additional Global Commands that are available throughout all ASN-9000 subsystems.

**Chapter 6 - System Subsystem Commands** - Describes commands for controlling session environment and saving those commands in a file for use in other sessions.

**Chapter 7 - NVRAM Subsystem Commands** - Describes commands for controlling session environment and saving those commands in a file for use in other sessions.

**Chapter 8 - Media Subsystem Commands** - Describes how the ASN-9000 relates to the physical media and bridging configuration information.

**Chapter 9 - Host Subsystem Commands** - Describes commands for modifying the TELNET and TCP configuration used by the ASN-9000 when using a TELNET connection to configure and manage the ASN-9000.

**Chapter 10 - TFTP Subsystem Commands** - Describes commands for performing file transfers between the ASN-9000 and TFTP servers.

**Chapter 11 - Bridge Subsystem Commands** - Describes commands for customizing the ASN-9000 bridge configuration.

**Chapter 12 - SNMP Subsystem Commands** - Describes commands for customizing SNMP communities.

**Chapter 13 - Internet Protocol (IP) Subsystem Commands** - Describes commands for configuring the ASN-9000 as an IP router.

**Chapter 14 - IP Multicast Subsystem Commands** - Describes commands for configuring the ASN-9000 for IP Multicast routing.

**Chapter 15 - IP/RIP Subsystem Commands** - Describes commands for configuring the ASN-9000 to exchange IP route information using Routing Information protocol (RIP).

**Chapter 16 - IP/OSPF Subsystem Commands** - Describes commands for configuring the ASN-9000 to exchange IP route information using Open Shortest Path First (OSPF) protocol.

**Chapter 17 - AppleTalk (ATALK) Subsystem Commands** - Describes commands for configuring the ASN-9000 as an AppleTalk router.

**Chapter 18 - IPX Subsystem Commands** - Describes commands for configuring the ASN-9000 as an IPX router.

**Chapter 19 - Configuring IPX Translation Bridging** - Describes commands for configuring the ASN-9000 for IPX encapsulation bridging.

**Chapter 20 - DECnet Subsystem Commands** - Describes commands for configuring the ASN-9000 as a DECnet router.

**Appendix A - Configuration Defaults** - Lists the factory defaults for all the ASN-9000 configuration parameters.

**Appendix B - Well-Known Ports** - Provides a pointer to RFC 1340, the “Well-known Ports” RFC.

## Related Publications

---

The following manuals are referenced throughout this manual. These manuals, including this manual, comprise the *ForeRunner* ASN-9000 Reference Manual set.

- *ForeRunner* ASN-9000 Hardware Reference Manual, MANU0255-01, November 7, 1997
- *ForeRunner* ASN-9000 Filters Reference Manual, MANU0280-01, November 7, 1997
- *ForeRunner* ASN-9000 ATM Software Reference Manual, MANU0273-01, November 7, 1997

## Technical Support

---

In the U.S.A., you can contact FORE Systems' Technical Support using any one of the following methods:

1. If you have access to the Internet, you may contact FORE Systems' Technical Support via e-mail at:

**support@fore.com**

2. You may FAX your questions to "support" at:

**412-742-7900**

3. You may send questions, via U.S. Mail, to:

**FORE Systems, Inc.  
1000 FORE Drive  
Warrendale, PA 15086-7502**

4. You may telephone your questions to "support" at:

**800-671-FORE or 412-635-3700**

Technical support for non-U.S.A. customers should be handled through your local distributor.

No matter which method is used for support, please be prepared to provide your support contract ID number, the serial number(s) of the product(s), and as much information as possible describing your problem/question.

## Typographical Styles

---

Throughout this manual, all specific commands meant to be entered by the user appear on a separate line in bold typeface. In addition, use of the Enter or Return key is represented as <ENTER>. The following example demonstrates this convention:

```
cd /usr <ENTER>
```

File names that appear within the text of this manual are represented in the following style: "...the `fore_install` program installs this distribution."

Command names that appear within the text of this manual are represented in the following style: "...using the `flush-cache` command clears the bridge cache."

Subsystem names that appear within the text of this manual are represented in the following style: "...to access the `bridge` subsystem..."

Parameter names that appear within the text of this manual are represented in the following style: "...using `<seg-list>` allows you to specify the segments for which you want to display the specified bridge statistics."

Any messages that appear on the screen during software installation and network interface administration are shown in `Courier` font to distinguish them from the rest of the text as follows:

```
.... Are all four conditions true?
```

## Important Information Indicators

---

To call your attention to safety and otherwise important information that must be reviewed to ensure correct and complete installation, as well as to avoid damage to the FORE Systems product or to your system, FORE Systems utilizes the following **WARNING/CAUTION/NOTE** indicators.

**WARNING** statements contain information that is critical to the safety of the operator and/or the system. Do not proceed beyond a **WARNING** statement until the indicated conditions are fully understood or met. This information could prevent serious injury to the operator, damage to the FORE Systems product, the system, or currently loaded software, and is indicated as follows:

### **WARNING!**



Hazardous voltages are present. To reduce the risk of electrical shock and danger to personal health, follow the instructions carefully.

**CAUTION** statements contain information that is important for proper installation/operation. Compliance with **CAUTION** statements can prevent possible equipment damage and/or loss of data and are indicated as follows:

### **CAUTION**



You risk damaging your equipment and/or software if you do not follow these instructions.

**NOTE** statements contain information that has been found important enough to be called to the special attention of the operator and is set off from the text as follows:



If you change the value of the LECS control parameters while the LECS process is running, the new values do not take effect until the LECS process is stopped, and then restarted.

## Laser Notice

---

**Class 1 Laser Product:**  
**This product conforms to**  
**applicable requirements of**  
**21 CFR 1040 at the date of**  
**manufacture.**

Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation.



The Laser Notice section applies only to products or components containing Class 1 lasers.

## Safety Precautions

---

For your protection, observe the following safety precautions when setting up equipment:

- Follow all warnings and instructions marked on the equipment.
- Ensure that the voltage and frequency of your power source matches the voltage and frequency inscribed on the equipment's electrical rating label.
- Never push objects of any kind through openings in the equipment. Dangerous voltages may be present. Conductive foreign objects could produce a short circuit that could cause fire, electric shock, or damage to your equipment.

## Modifications to Equipment

Do not make mechanical or electrical modifications to the equipment. FORE Systems, Inc., is not responsible for regulatory compliance of a modified FORE product.

## Placement of a FORE Systems Product

### CAUTION



To ensure reliable operation of your FORE Systems product and to protect it from overheating, openings in the equipment must not be blocked or covered. A FORE Systems product should never be placed near a radiator or heat register.

## **Power Cord Connection**

### ***WARNING!***



FORE Systems products are designed to work with single-phase power systems having a grounded neutral conductor. To reduce the risk of electrical shock, do not plug FORE Systems products into any other type of power system. Contact your facilities manager or a qualified electrician if you are not sure what type of power is supplied to your building.

### ***WARNING!***



Your FORE Systems product is shipped with a grounding type (3-wire) power cord. To reduce the risk of electric shock, always plug the cord into a grounded power outlet.

# CHAPTER 1

## Features Overview

This chapter provides an overview of the major features of the ASN-9000. Specifically, this chapter discusses:

- Intelligent Packet Switching
- Software
- Network Management

### 1.1 Intelligent Packet Switching

---

Much of the packet switching in the ASN-9000 is performed by the ASN-9000 Packet Engine. The ASN-9000 Packet Engine (PE), is the centralized packet processing and forwarding engine of the ASN-9000. When a packet is received on a segment attached to the ASN-9000, the packet is forwarded to the PE and placed in Shared Memory, where it is examined and either dropped or forwarded, as applicable. Additionally, packet switching is also accomplished within the ATM Network Interface Modules (NIMs), for packets destined for connections on the same NIM. The ASN-9000 supports the ASN-9000 PowerCell NIM. The PowerCell NIM can support two ATM Media Adapters (AMAs), one as a primary ATM physical (PHY) interface and the other as a backup. There are three variations of AMAs that are supported by the ASN-9000. These include OC-3 Single-Mode, OC-3 Multimode and OC-3 UTP. Refer to the *ForeRunner ASN-9000 Hardware Reference Manual* for detailed information on the Intelligent Packet Switching features of the ASN-9000, the ASN-9000 PowerCell and supported AMAs.

### 1.2 Software Features

---

The following sections describe the features available in the ASN-9000 software. These features include:

- System Management
- Virtual Local Area Networks (VLANs)
- Bridging and Routing

## 1.2.1 System Management Features

The system management features discussed in this section focus on the management of the ASN-9000, rather than configuration and management of network interfaces. Some of the items discussed include:

- Multiprocessor Optimization
- Multiple Boot Sources
- UNIX-Like Command Line Interface
- Local File Management System
- Concurrent Command Line Sessions
- System Configuration Files
- Session Parameter Files
- Automatic Segment State Detection
- Segment Statistics

### 1.2.1.1 Multiprocessor Optimization

The goal of Multiprocessor Optimization is to minimize the latency caused in the normal packet forwarding functions due to the processing of management events. By moving these processing-intensive functions to a separate MCPU, the latency of packets in the fast path can be kept to a minimum.

This feature is dependent on having a ASN-9000 Packet Engine with a Packet Accelerator installed. With the accelerator installed, there are four CPUs available. Without the Multiprocessor Optimization feature, only three of these CPUs are used. This feature makes use of the fourth CPU by splitting the functions of the single MCPU in the current ASN-9000 software architecture.

Multiprocessor Optimization moves all of the fast path packet processing to one MCPU that retains the slow path and management functions on the other MCPU. Multiprocessor Optimization automatically detects the presence of an Accelerator Card at boot time and operates in the appropriate mode. Without the Accelerator Card, the system uses only one MCPU for all functions.

### 1.2.1.2 Multiple Boot Sources

The ASN-9000 software can be configured to boot from any one, or combination, of the following three boot sources: floppy drive, Flash Memory Module, or TFTP/BOOTP file server. The boot order can be configured in NVRAM to ensure against primary boot source failure.

### 1.2.1.3 UNIX-Like Command-Line Interface

The ASN-9000 is managed through the use of UNIX-like commands. These commands are issued from a management terminal, either directly through a TTY connection or indirectly through an in-band (TELNET) connection.

### 1.2.1.4 Local File-Management System

The ASN-9000 operating system contains global commands that allow files to be displayed, copied, renamed or removed from either the floppy disk or Flash Memory Module. Additionally, file checksums can be calculated, display directory and volume information for the floppy drive or Flash Memory Module, and if necessary, reformat the Flash Memory Module.

### 1.2.1.5 Concurrent Command-Line Sessions

Up to four management sessions can be open at the same time. The primary session is always on TTY1. A second TTY session can be opened on TTY2. Additionally, up to two TELNET sessions can also be open, all simultaneously.

### 1.2.1.6 System Configuration Files

This section tells users they can preserve configuration changes they effect through software commands by saving the changes in a ASN-9000 configuration file. If they save the changes to the file name `cfg`, the changes are automatically applied to the switch following a software reboot, provided the `cfg` file is present on the boot source used to boot the switch.

### 1.2.1.7 Session Parameter Files

The ASN-9000 software contains commands to modify parameters that control user sessions. These parameters include scroll control, TELNET control characters, commands aliases, and timed commands. Changes to the defaults for these session parameters are lost when the session is closed.

Changes to the session parameters can be saved in an environment file. At any time during a user session, an environment file can be read (load) and reinstate the session parameter changes stored in the file.

If an environment file is saved under the name `root.env`, the file is automatically loaded whenever logging on with root status. Likewise, saving an environment file under the name `monitor.env`, the environment parameters in that file are automatically loaded when logging on with monitor status or changing the user level from root to monitor during a session.

### 1.2.1.8 Automatic Segment-State Detection

Automatic Segment-State Detection, when enabled, automatically senses when a link (or something configured on the link) is “bad” or “down”. When a “bad” or “down” link is detected on a particular port, the state of the segment is reflected in the software’s interface tables. *ForeView* Network Management software allows link types to be enabled or disabled on a particular port.

### 1.2.1.9 Segment Statistics

Access method and protocol statistics related to segment and packet activity on the ASN-9000 being managed can be viewed. For example, state-change statistics for individual segments can be displayed to see how many times a particular segment has gone up or down since the software was last booted. Statistics related to the protocols are briefly described in Section 1.2.3.6.

## 1.2.2 Virtual Local Area Networks (VLANs)

A VLAN is a collection of segments that share the same group name or protocol interface address. Bridging segments are a Layer-2 VLAN.

A Layer-2 VLAN are created when creating a bridge group. The ASN-9000 comes with a default bridge group, called `default`, that contains all the ASN-9000 segments.

A Layer-3 VLAN is created by assigning the same IP, IPX, or AppleTalk interface address to multiple segments. When the ASN-9000 determines a packet is to be sent to a Layer-3 VLAN assigned to multiple segments, the software forwards a copy of the packet on each segment. When this happens, from a physical standpoint, a separate packet has been sent out each physical interface; however, from a logical standpoint, the forwarded packet has been forwarded onto its single destination network or subnet, irrespective of how many physical interfaces that network or subnet is configured on.

## 1.2.3 Bridging and Routing Features

The `bridge` subsystem contains commands for configuring and managing the ASN-9000 as an IEEE 802.1d bridge. Up to 32 network (bridge) groups can be defined, each containing any subset of segments.

### 1.2.3.1 Bridge Table and Bridge Cache

The software maintains a bridge table containing the MAC-layer hardware addresses of devices to which the ASN-9000 is able to bridge packets. The software maintains the table by automatically adding new entries and deleting unused entries. In addition, individual entries can be added or removed, including entries supporting multi-homed hosts.

Here is an example of the bridge table. Although only a handful of bridge entries are shown in this example, the bridge table usually contains many entries.

```
Bridging table (aging time = 60 minutes)
Ethernet-address  Seg Rule  Flags
00-00-00-00-00-00 01 none   aged
00-00-c0-ea-9f-17 01 none
08-00-20-10-19-ac 08 none
00-00-c0-ed-61-4a 01 none
08-00-20-0c-5a-48 08 none
02-cf-1f-90-40-23 01 none
08-00-20-0c-3a-a2 02 none
08-00-20-0c-5a-d2 08 none   aged
ff-ff-ff-ff-ff-ff --1 none   permanent  bmcast
```

In addition to the bridge table, the ASN-9000 maintains a *bridge cache* of the most recently used source-destination pairs. A *source-destination pair* contains a packet's source MAC-address and destination MAC-address. The cache provides a fast path for the bridging software and provides an at-a-glance view of current bridging activity. The bridge cache can be displayed to see at-a-glance the source-destination pairs that are being used frequently.

#### 1.2.3.1.1 802.1d

The ASN-9000 switch can be used “right out of the box” as an 802.1d Bridge. The designation 802.1d refers to the IEEE committee number that came up with the spec for this type of bridge. For more information regarding 802.1d bridging, refer to RFCs, 1493 and 1525.

#### 1.2.3.1.2 Spanning-Tree

The ASN-9000 bridge software includes implementation of the 802.1d Spanning-Tree algorithm. When enabled, the algorithm on the ASN-9000 identifies and “breaks” loops in the network, without requiring configuration changes. Commands in the `bridge` subsystem can be used to fine-tune Spanning-Tree parameters to fit the network.

#### 1.2.3.1.3 IPX Translation Bridging

IPX translation bridging lets the user configure one or more IPX networks that span across Ethernet segments using different packet encapsulations. This type of bridging is different from 802.1d bridging, which bridges packets based on the MAC-layer hardware address of the devices in the network. IPX translation bridging is used only in IPX networks.

### 1.2.3.2 IP Routing

Commands in the `ip` subsystem are available to configure ASN-9000 segments for IP routing. Using `ip` commands, IP interfaces can be assigned to individual segments. The IP routing software also supports IP VLANs, enabling a single IP subnet to be defined that spans multiple ASN-9000 segments. The following subsections describe major features of the `ip` subsystem. Refer to *Chapter 13, IP Subsystem Commands* in this manual for more information about these features and the `ip` commands.

#### 1.2.3.2.1 Routing Information Protocol (RIP)

The `ip/rip` subsystem commands enable the ASN-9000 to perform IP routing. Using commands in this subsystem, RIP parameters such as `talk` and `listen` can be configured and statistics displayed on a segment-by-segment basis.

#### 1.2.3.2.2 Open Shortest Path First (OSPF)

The `ip/ospf` subsystem contains commands to configure the ASN-9000 as an OSPF router. OSPF is a routing protocol that enables each participating router to use a topological map of the network to route packets. OSPF routers exchange route information using link-state advertisements (LSAs). An LSA is a packet that reports link states (up or down) of a router's interfaces that are attached to devices in the OSPF network.

#### 1.2.3.3 AppleTalk Routing

The `atalk` subsystem contains commands to configure ASN-9000 segments for AppleTalk Phase-2 routing. AppleTalk zones and interfaces can be defined, and ping AppleTalk nodes.

#### 1.2.3.4 IPX Routing

The ASN-9000 can be configured as an IPX router. Additionally, the ASN-9000 provides management information about IPX routes and servers through implementation of IPX Routing Information Protocol (RIP) and Service Advertisement Protocol (SAP). RIP or SAP `talk` and `listen` parameters can be selectively enabled on a per-segment basis to control the flow of RIP and SAP updates in the network.

#### 1.2.3.5 DECnet Routing

The `dec` subsystem contains commands for configuring the ASN-9000 to perform DECnet Phase IV routing. Depending upon the network configuration, the ASN-9000 can be configured to function as a Level-1 or Level-2 router. DECnet statistics for the ASN-9000 can be displayed (in its capacity as a DECnet node) for the individual segments configured as DECnet interfaces.

### 1.2.3.6 Route Protocol Statistics

Statistics can be gathered for the following Internet routing protocols:

- AppleTalk
- Bridge
- DECnet
- IP
- IPM
- IPX
- OSPFv2
- RIP
- SNMP
- TCP
- TCP/IP

### 1.2.3.7 Security Filters

Filters can be defined and applied to segments or protocol interfaces to control the traffic sent and received on the segments or interfaces. The following types of filters can be defined:

- Bridge filters
- Host (TCP) filters
- IP filters
- IP route filters (RIP and OSPF)
- AppleTalk filters
- IPX RIP and SAP filters

## 1.3 Network Management Features

---

The ASN-9000 has a rich management environment providing comprehensive support for SNMP, as well as local RS-232 and Telnet console support. *ForeView* graphical network management software provides true point-and-click device configuration and runs on a variety of popular management stations.

### 1.3.1 Network Management System (NMS)

The NMS manages the ASN-9000 by sending a request to a software module, or agent, to change the value of one or more variables on the device. For example, an agent reports data such as the number of incoming and sent packets, or the number of dropped packets on that

device. Then, the managed device and the NMS use Simple Network Management Protocol (SNMP) as the common protocol language to exchange the information requested by the NMS.

### 1.3.2 Management Information Base (MIB) Agents

MIBs contain the definitions of all the resources (represented by managed objects within the MIB) that are managed by a network management system (NMS). The managed object has properties that hold values, such as ASN-9000 routing table information, error counters, and so on.

### 1.3.3 ForeView

ForeView is a graphical-based management application providing a simplified tool for managing the ASN-9000. With a point-and-click interface, *ForeView* provides access to ASN-9000 functions, both system-level and segment-level. *ForeView* can monitor errors, control ASN-9000, bridge, and routing configuration parameters, and display, print, and save statistics.

*ForeView* software integrates the ASN-9000, bridge, and router features into a single application with access and control of all information from one location. The software also features fault management to troubleshoot, analyze, and monitor multiple Ethernet segments using a single network analyzer.

ASN-9000 statistics are shown in graphical formats, and the physical attributes of the managed ASN-9000, such as model and segment type, are displayed on the front panel of a graphical representation of the ASN-9000 being managing. The graphical representation is displayed when starting *ForeView*. For more information about the *ForeView* Network Management application, see *ForeView Network Management User's Manual*.

This chapter describes the different types of software used by the ASN-9000, and describes the files on the system software diskettes shipped with the ASN-9000. Upgrade the software in the following order:

- 1. Packet Engine boot PROM.
- 2. Boot PROMs on intelligent NIMs. Software files.
- 3. Runtime PROM Images to Flash Memory.

For information on upgrading firmware, see the following:

- Contact FORE Systems TAC for procedures to upgrade the Packet Engine boot PROM.
- Refer to the *ForeRunner ASN-9000 Release Notes* for procedures to upgrade the runtime PROMs on intelligent NIMs.

## 2.1 Types

The ASN-9000 modules use the following types of software:

<b>Packet Engine boot PROM</b>	Contains software used by the Packet Engine when it is booted. From this PROM, configuration values can be changed—including the boot source—stored in NVRAM. The command prompt for the Packet Engine boot PROM looks like this:  <PROM-7PE>
<b>System software</b>	Sometimes called “runtime software.” Contains most of the commands and features documented in this manual. The runtime software can be accessed from the runtime command prompt. The runtime prompt appears as:  1:ASN-9000:
<b>NIM Boot PROM</b>	On intelligent NIMs (such as 6x1FE, and PowerCell ATM modules), the boot PROM contains software used by the module when it is booted.

<b>Runtime PROM</b>	Contains runtime features used by intelligent NIMs. The runtime PROM software for each intelligent NIM is stored in firmware on the NIMs themselves.
---------------------	--

## 2.1.1 Software Files

Each ASN-9000 is shipped with two sets of 3.5-inch high-density (HD) floppy diskettes that contain the latest version of the system software, boot PROMs runtime PROMs for the intelligent NIMs, and a boot definition file. The two sets of diskettes are identical and each contains the following files:

7PE	System software image file; this file resides on the boot source and is loaded when the system is initially boot.
ppu-7PE	Packet-Engine boot PROM image file; this file is resident on the boot PROM.
7atm	Runtime PROM image for ASN-9000 PowerCell 700.

## 2.2 Other Files Supplied with the ASN-9000

---

Additional files sent with the ASN-9000 are files used for testing the system.

bootdef	Used by the system when it boots the software to identify the name of the system software image file and the configuration file.
dispcfg	Configuration file that runs a series of commands that display system configuration information and statistics; this file is useful for diagnosing configuration problems.

## 2.2.1 Files the ASN-9000 Creates

In addition to the files shipped with the system, the following types of files can be created and saved onto a software diskette or Flash Memory Module during ASN-9000 operation:

<b>cfg</b>	ASN-9000 configuration file; created when the <b>savecfg cfg</b> command is issued. Configuration files can be saved under any other DOS-compatible filename, but the configuration must be manually loaded, unless the user also edits the <b>bootdef</b> file to contain the configuration file name.
<b>root.env</b>	Environment file for root sessions; created when the <b>saveenv root.env</b> command is issued. Environment files can be saved under any DOS-compatible filename, but the file must be manually loaded.
<b>monitor.env</b>	Environment file for monitor sessions; created when the <b>monitor.env</b> command is issued.
<b>ASN-9000.dmp</b>	Dump file; if the ASN-9000 system crashes, a dump file is created on a local storage device.
<b>iop1.dmp</b>	Another dump file that the software can produce when a crash happens.
<b>iop2.dmp</b>	Another dump file that the software can produce when a crash happens.



# CHAPTER 3

## Using the Command Line Interface

This chapter describes how to use the command-line interfaces. The chapter will contain the following sections:

- Using the new user interface
- Displaying on-line help for the new user interface
- Issuing commands for the new user interface

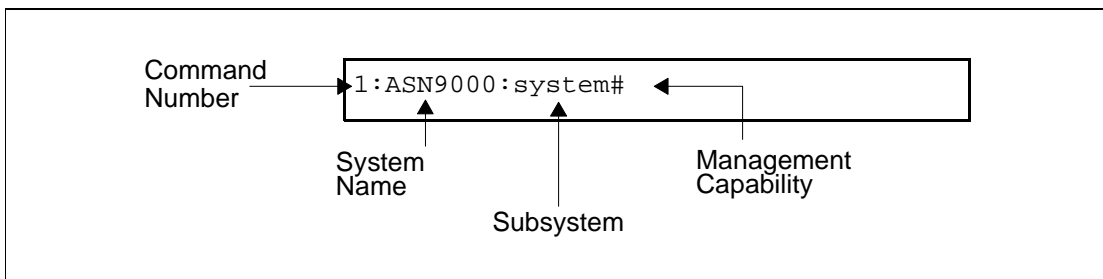
### 3.1 Using the New User Interface

---

This section describe how to use the new user interface in the following sections.

#### 3.1.1 Issuing Commands

Regardless of whether you are accessing the *ForeRunner* ASN-9000 through a TTY (RS-232) port or through a TELNET session, you issue commands at a command prompt such as the one shown in Figure 3.1:



**Figure 3.1 - User Interface Command Line**

As shown in Figure 3.1, the command prompt has four components:

<b>Command Number</b>	<p>The sequential number of the command in this session (similar to a command number in the UNIX C-shell). In this example, the command number is 1.</p> <p>When you issue a carriage return, the <i>ForeRunner</i> ASN-9000 attempts to execute the command you have entered at the command prompt. A message or data (if requested) is displayed, then a new command prompt is displayed. The number in the command prompt is one number higher than the number in the previous command prompt.</p>
<b>System Name</b>	<p>The name assigned to this ForeRunner ASN-9000. The default name, ASN-9000, can be changed using the <b>system sysname</b> command. In this example, the name is ASN9000.</p>
<b>Subsystem</b>	<p>The name of the subsystem currently in use. Commands issued at the command prompt must either be global commands or commands within the current subsystem. In this example, the subsystem is <i>main</i>.</p>
<b>Management Capability</b>	<p>Indicates whether the session is in <i>monitor</i> capability or <i>root</i> capability:</p> <p>&gt; Indicates monitor (display only) capability. <i>Monitor</i> capability lets you display statistics and read configuration information. You cannot issue commands that change the configuration, clear statistics, or modify internal tables.</p> <p># Indicates root (configuration) capability. <i>Root</i> capability lets you issue any command, including commands that change the configuration and the internal tables.</p> <p>In this example, the management capability is shown as # for root capability.</p>



If you start a user session and the `login:` prompt is displayed, rather than the command prompt, you must enter the appropriate password before proceeding.

The command prompt described in this section lets you issue system software commands. The `<PROM-7PE>` prompt also can be used to issue some commands, including NVRAM commands, but does not allow most system software commands.

### 3.1.2 Entering and Editing Command Text

All commands are typed at the command prompt using a keyboard attached to the workstation, terminal, or PC you are using as a management station. The workstation or terminal must be attached to one of the TTY ports or connected to the switch through a TELNET session.

Commands and arguments are case-sensitive and should be entered only as shown in the manual or on-line help. Each command must fit on a single line and cannot exceed 128 characters in length. The keys you use to edit and issue commands are the standard keys used on most UNIX workstations:

- To issue a command, enter the command name and arguments (if needed) after the command prompt, then press `<Enter>`.
- To erase individual characters in a command, use the `<Backspace>` or `<Delete>` key, or the `EraseChar` character assigned in your TELNET session (usually `<Ctrl+H>`).
- To cancel an entire line of input, use the reassign character (usually `<Ctrl+U>`).
- To control the scrolling of output on the terminal, use `<Ctrl+S>` to stop the flow and `<Ctrl+Q>` to resume the flow.

Commands in the `host` subsystem can be used to display or change the key sequences used in TELNET sessions. Key sequences can be changed or displayed for the current session or the default key sequences used for all sessions.

### 3.1.3 Displaying a Command's Syntax

Entering a command without parameters from the boot PROM or runtime prompts results in the system responding with a display of the proper syntax to be used for the command, unless there are no parameters for the command entered. Additionally, the `help|h|?` command offers three levels of help.

The first level of on-line help gives a command's syntax based on the search of a *verb* command (action) and describes the command as in the following example:

```
8:ASN-9000:ip# dd
[show] dd|default-device
```

The second level of on-line help gives a command's syntax based on the search of a *noun* command (object) and describes the command. In the following example, the `ip` subsystem is used to display a portion of the help available for the `arp` command:

```
9:ASN-9000:ip# help arp show
arp [show] [-r] [-t] [-s] [<disp-restrictors>]
```

By specifying additional parameters to a command, more specific on-line help is available as in the following example:

```
10:ASN-9000:ip# help arp show age
arp [show] age Show ARP aging time
```

## 3.1.4 Command Syntax

This section describes and gives examples for the new user interface syntax. The following subsections describe the syntactical elements of the new user interface:

### 3.1.4.1 Verbs

This section will describe the verbs in the new user interface in the following subsections.

#### 3.1.4.1.1 set and unset

These verbs set or remove the setting from system parameters. Examples include:

- The boot order (the order in which the switch tries to use the floppy drive, Flash Memory Module, or TFTP boot server for booting the software)
- stty (scroll control) parameters
- Timed commands
- Routing protocols
- Specific bridging and routing protocol features

When the **set** or **unset** verb is prepended by **c** or **p**, or **n** or **s**, the verb applies only to specific segments (**c** or **p**) or specific networks (**n** or **s**).

#### 3.1.4.1.2 define and undefine

These verbs create or delete templates and rules, both of which are components of filters.

#### 3.1.4.1.3 attach and detach

These verbs apply or remove filters (created using the **define** and **undefine** verbs) to *ForeRunner* ASN-9000 segments or protocol interfaces.

#### 3.1.4.1.4 add and delete

These verbs add or delete objects from tables. Examples include bridge-table entries, protocol interface-table entries (IP, AppleTalk, IPX, and DECnet), and route-table entries.

#### 3.1.4.1.5 enable and disable

These verbs turn on or off specific software features. Examples include bridging and protocol routing, specific IP routes, IP Helper, and so on.

When the **enable** or **disable** verb is prepended by **p**, or **n** or **s**, the verb applies only to specific segments (**p,s**) or specific networks (**n** or **s**).

#### 3.1.4.1.6 show and clear

These verbs display and clear (if applicable), configuration information, tables, caches, and statistics. With these verbs you also can display configuration information or display and clear statistics.

### 3.1.4.2 Objects

This section will describe the verbs in the new user interface in the following subsections.

#### 3.1.4.2.1 config

This object shows the parameter settings for the switch hardware, the bridging subsystem, and the routing protocol subsystems. In general, the **config** (or **show config**) command displays system parameters that can be configured through the software.

#### 3.1.4.2.2 status

This object shows the current status for the switch hardware (segment up/down status, 10Base-T port status, and so on), the current bridge status of *ForeRunner* ASN-9000 segments (bridging enabled or disabled, Spanning Tree enabled or disabled), and so on. In general, the configuration parameters are displayed when you issue the **status** (or **show status**) command.

#### 3.1.4.2.3 stats

When the **stats** command (or **show stats**) command is issued, the statistics related to the feature area in the current subsystem (or specified subsystem, if different from the current) are displayed. For example, the **stats** command issued from within the **ip** subsystem displays IP, ARP, and ICMP packet statistics.

#### 3.1.4.2.4 interface

This object is a routing protocol interface (IP, IP Multicast, AppleTalk, or IPX).

#### 3.1.4.2.5 route

This object is an IP, IP Multicast, AppleTalk, IPX, or DECnet route.

#### 3.1.4.2.6 bt

This object is the bridge table.

#### **3.1.4.2.7 cache**

This object contains “fast-path” entries. The bridge subsystem and all the routing protocol subsystems contain caches. The “fast-path” entry is a shortcut the *ForeRunner* ASN-9000 software uses to bridge or route packets. When a bridge table or route table is in the fast path, the *ForeRunner* ASN-9000 software does not need to perform all the bridging or routing processing that it normally performs in order to bridge or route a packet. The *ForeRunner* ASN-9000 software maintains the caches by placing in them the most recently used source destination MAC-address pairs (for bridging) or protocol interface addresses (for routing).

### **3.1.4.3 Parameters**

This section will describe the types of parameters in the new user interface.

#### **3.1.4.3.1 Keyword Parameters**

In the new user interface, a keyword parameter is a parameter that you can enter at any point following the verb.

#### **3.1.4.3.2 Positional Parameters**

A positional parameter is a parameter that you must enter in a specific position following the verb of a syntactical command. The need for positional parameters in the new user interface is infrequent because the software uses keywords to determine the function you’re performing on the *ForeRunner* ASN-9000. When the need for a positional parameters arises, the software let’s you know and provides you with the correct position of parameters in the command.

# CHAPTER 4

## Global Commands

This chapter describes how to use the global commands to perform the following tasks:

- Boot the software
- Log in and log out
- Change management capability between root and monitor
- Control scroll
- Display and change command-line control characters (sometimes referred to as TELNET control characters, but they apply to TTY sessions too)



There are global commands that let you do other things (display help, list subsystems), which are described in *Chapter 5, More Global Commands*.

### 4.1 Rebooting the Switch

---

You can boot (or reboot) the system software using any of the following methods:

- Press the reset switch (labeled RST), located on the front of the Packet Engine.
- Issue the `system reboot` command.
- Issue the `boot (b)` command at the `<PROM-7PE>` prompt.
- Turn the power supplies off, then back on.

#### 4.1.1 Reset Switch

The reset switch is located to the right of the floppy drive and is labeled RST. When you press the reset switch, the Packet Engine performs a “cold” restart of the ASN-9000. During a cold restart, the Packet Engine conducts a power-on self-test to check its various hardware components.

Depending on the boot preference(s) you specify, the Packet Engine uses files on a floppy disk, Flash Memory Module, or on a TFTP server (network booting) to configure the ASN-9000 for runtime operation. For an example of the boot messages displayed by the ASN-9000 and additional methods for rebooting, see the *ForeRunner ASN-9000 Hardware Reference Manual*.

## 4.1.2 Using the Reboot Command

Issue the **reboot** command from the **system** subsystem to reboot the ASN-9000. The following is an example of the boot messages that are displayed.

```
69:ASN-9000:system# reboot
FORE 7000 PE
Prom version: 7pep-2.5.3-Alpha2 (s1.80) 1996.02.13 14:21
I-cache 16K OK
Entering cached code
I-cache execution OK
D-cache 4K OK
SRAM 128K OK
DRAM .....24064K OK
Shared Memory ...2048K OK
Entering Monitor
FlashInit: found 2MB Flash Memory Module
Board Type: 7PE , CpuType: MCPU, Instance: 2
Ethernet address: 00-00-ef-02-b9-c0
(normal start)
Hit any key now to abort boot [0]:
Trying floppy boot...
Boot definition file: bootdef (default)
Using disk bootdef, parsed as version 0
Loading file "7pe" (AB format) |
File loading complete
initial program counter: 0x80210b40
Disabling break interrupt 4
Switching interrupt source back to ID bits
FORE ForeRunner ASN-9000 Runtime System Software
7pe-BRULEE-Alpha33 (s1.445) 1996.12.05 00:45
tty driver initialized
system timer initialized
floppy driver initialized
FlashInit: found 2MB Flash Memory Module
memory manager initialized
Board Type: 7PE , CpuType: MCPU, Instance: 2
System ethernet address: 00-00-ef-02-b9-c0
Looking for packet accelerator card
Did not find packet accelerator - will use 2MB shared memory
Shared mem available: 2024448
PE: slot 5
X Bus:
Slots that are equipped:
Slots that are equipped and latched:
Y Bus:
Slots that are equipped: 2 1
Slots that are equipped and latched: 2 1
```

The boot messages shown in this example are displayed when the software is booted from the floppy drive. You can configure the ASN-9000 to boot the software from the Flash Memory Module (if installed) or from a TFTP file server.


**NOTE**

If a NIM fails to come up when you reboot, make sure the left ejector handle is pressing on the activation switch, located behind the ejector handle. If the switch is not being pressed, the NIM does not operate.

### 4.1.3 Booting the ASN-9000

Use the **boot** command to boot the ASN-9000 from the <PROM-7PE> prompt. Here is the syntax for this command:

```
boot|b -n [fd|fm]
```

**fd** Boots the software from the floppy drive.

**fm** Boots the software from the Flash Memory Module.

If you do not specify a boot source, the boot order configured in NVRAM is used. If you have not configured a boot order, the floppy drive (**fd**) is used.


**NOTE**

The system software image file (7PE) must already be present on the boot source you specify. The **boot** command does not affect the boot order specified in NVRAM.

## 4.2 Logging In and Out

---

This section describes how to log in to the switch through a direct connection (TTY) or an in-band connection (TELNET). This section also describes how to log out.

## 4.2.1 Logging In

When the Lock Switch is set to the unlocked position (U), you do not need to log in. When you boot (or reboot) the ASN-9000, a command prompt is displayed. However, if the Lock Switch is set to the Locked position (L), the login: prompt is displayed. You must enter “root” or “monitor,” then enter a password. The password you enter depends upon the management capability you want:

- If you specify “monitor,” enter the password for monitor capability.
- If you specify “root,” enter the root password.

When the ASN-9000 is shipped from the factory, the password for each management capability is blank. At the password: prompt, press <Enter>. To set or change a password, use the **system passwd** command.

### 4.2.1.1 TTY

The ASN-9000 comes with two external RS-232 (TTY) ports that can be used to connect management terminals to the PE2. The primary management session is always through TTY1, so TTY1 is the default used when booting the ASN-9000. When you boot the ASN-9000, a command prompt will appear on the management terminal attached to TTY1. For TTY2, issue the following command:

```
system enable tty2
```

#### 4.2.1.1.1 Baud Rates

When the Lock Switch is unlocked, the ASN-9000 always uses the default baud rates for its TTY ports: 9600 for TTY1 and 1200 for TTY2. Each ASN-9000 port supports baud rates of 1200, 2400, 4800, 9600, 19200, 57600, and 115,200. If you need to use a different baud rate for either port, you must use the **system setbaud** command to specify the baud rate you need.

When you set the baud rate using the **system setbaud** command, the rate is recorded in NVRAM. However, the baud rates recorded in NVRAM are not used unless the Lock Switch is on when you boot the ASN-9000. If the Lock Switch is off, the baud rates stored in NVRAM are ignored and the default rates are used instead.

See the *ForeRunner ASN-9000 Hardware Reference Manual* for instructions on setting the baud rate for a TTY port.

NOTE

Although you can change the baud rate on TTY1, you cannot do so until you log in and issue the **system setbaud** command. The first time you access the user interface on the ASN-9000, the management terminal or modem attached to the TTY1 port must be set to 9600 baud.

#### 4.2.1.2 TELNET

Configure an IP interface on the segment attaching the management terminal to the switch, then issue the appropriate command from the workstation's command prompt (not an ASN-9000 prompt):

```
telnet <switch's IP address>
```

#### 4.2.2 Logging Out

To log out from a ASN-9000 session, issue one of the following commands:

- **logout**
- **bye**

These commands end the session from which the command is issued, but do not affect other user sessions on the ASN-9000. You can end all sessions on the ASN-9000 by powering down the ASN-9000 or rebooting.

Configuration changes to the ASN-9000 remain in effect until the next software reboot. If changes are saved in a configuration file, they can be reinstated following reboots. To save environment changes, you must first save them to an environment file or the environment changes will be lost. Settings to the ASN-9000 can be reinstated by reading the environment file into a user session.

NOTE

Configuration changes you make affect only the switch. Environment changes affect the user session from which they are made but do not affect configuration of the switch.

## 4.3 Changing Management Capability

---

When the Lock Switch is set to the unlocked position (U), you do not need to log in. When you boot (or reboot) the ASN-9000, a command prompt is displayed. However, if the Lock Switch is set to the Locked position (L), the login: prompt is displayed. You must enter “root” or “monitor,” then enter a password. The password you enter depends upon the management capability you want:

- If you specify “monitor,” enter the password for monitor capability.
- If you specify “root,” enter the root password.

When the switch is shipped from the factory, the password for each management capability is blank. At the password: prompt, press <Enter>. To set or change a password, use the **system passwd** command.

This chapter describes how to perform the following tasks:

- Displaying and using command histories
- Managing files on the floppy drive and Flash Memory Module
- Defining and using command aliases
- Defining and using timed commands
- Saving and reading environment files (session settings)

### 5.1 Displaying and Using Command Histories

---

For each session, the *ForeRunner* ASN-9000 software maintains a history of the 32 most recently issued commands. Using the history commands, you can display the command history, reissue commands, or edit and reissue commands.

To display the 32 most recently issued commands, issue the **history** command.

To reissue or edit commands listed in the command history, use the *history control characters*.

Here are the default history control characters.

- ! History-prefix character.
- ^ Quick-substitution character.

You use the history control characters to form commands to reissue (or modify and reissue) commands from the command history. Here are the history commands you use to edit and reissue commands listed in the command history. The syntax is shown using the default history characters.

- !! Repeats the previous command.
- ! <n> Repeats a command listed in the command history, where <n> indicates the number of the command as listed in the history.
- ! <-i> Issues a previously issued command, where <i> is the offset back from the current command. For example, the command !-1 gives the same results as !!, reissuing the previous command.

<b>!<i>&lt;substring&gt;</i></b>	Repeats a previous command that begins with the string identified by <i>&lt;substring&gt;</i> .
<b>^<i>&lt;old&gt;</i>^<i>&lt;new&gt;</i></b>	Modifies, then reissues the previous command, where <i>&lt;old&gt;</i> indicates the string to be replaced with <i>&lt;new&gt;</i> .

Use the **histchars** command to display the current history control characters.

To change the history control characters, issue the **histchars** command with one or both optional arguments:

```
histchars [<ch1>[<ch2>]]
```

Here is an example of the use of this command:

```
5:ASN-9000:system# histchars  
history sub: !quick sub: ^
```

## 5.1.1 Resetting the Return Code

Use the **rcprompt enable** command to reset return codes for commands executed automatically from a script. Here is the syntax for the **rcprompt enable** command.

```
rcprompt enable|disable
```

<b>enable disable</b>	Enables printing of command return codes in the next UI prompt. Return codes will be displayed with 0 for successfully executed commands. This feature is intended primarily for automated interactions with the ASN-9000 command-line interface.
-----------------------	---

## 5.2 Managing Files

---

This section describes how to manage files on the local storage devices (Flash Memory Module and floppy drive) in the following sections:

### 5.2.1 Displaying a Directory

Use the **dir** command to display a directory for a single storage device or both devices:

```
ls|dir [<filespec>]
```

**<filespec>** Specifies a file name. You can use the wildcards (\*) and (?) for any portion of the file name.

This command is designed to present directory information in a way familiar to UNIX users. Here are some examples of the type of display produced by the **dir** command. In Figure 5.1, the contents of the Flash Memory Module are listed:

				File Size
File Name	61:ASN-9000:system# <b>dir fm:*. *</b>			
	FM:IP.cfg	242	5-05-1996	2:18p
	FM:RIP.cfg	107	7-12-1996	4:18p
	FM:PAUL.lbs	190	5-05-1996	2:18p
	FM:ALEX.lbs	210	5-05-1996	2:25p
	FM:TONY.lbs	180	5-05-1996	2:18p
	FM:LISA.lbs	165	1-05-1997	7:18p

Device on which file is stored

Time that file was written or changed on the device

Date that file was written or changed on the device

**Figure 5.1 - Details List of the Flash Memory Module**

## 5.2.2 Deleting a File

Use the **rm** to remove a file from a floppy disk or from the Flash Memory Module. Both commands do the same thing. Here is the syntax for the **rm** command:

```
rm [-f] [-i] <filespec> [<filespec>...]
```

- f** Forces the software to remove the file(s), without asking you if you are sure before removing each file.
- i** Overrides the **-f** (Force) flag, presenting a prompt before removing each file. The prompt gives you the opportunity to cancel your request to remove the file. If you do not specify **-f** or **-i**, **-i** is the default.

**<filespec>** Specifies the name of the file(s) you want to delete.

## 5.3 Using Command Aliases

---

The command-line interface provides an *alias* mechanism that lets you issue frequently-used commands with just a few keystrokes. Each time you need to issue the command, you can type the alias instead of the command itself. The ASN-9000 alias mechanism is a simplified version of the alias mechanism in the UNIX C-shell.

Aliases are local to the current command-line session. That is, they are not remembered across logins or resets of the ASN-9000 unless you save them to an environment file. Each session can have up to 32 aliases. Aliases can be stored in an environment file by issuing the **system saveenv** *<file-name>* command.

In general, when you issue an alias from the command line, it must be the first item after the command prompt. However, you can enter a subsystem name before the alias; for example:

```
media box
```

<b>media</b>	Is the subsystem name.
<b>box</b>	Is the alias.

You also can use an alias to make a frequently used command global; for example: **alias box media showcfg**.

In addition to entering an alias directly from the command line, you can use an alias as part of a timed command. When the timed command is activated, the command represented by the alias is issued.

### 5.3.1 Defining an Alias

To define an alias, issue the following command:

```
alias <name> <command>
```

<b>&lt;name&gt;</b>	Is the name that defines the alias.
<b>&lt;command&gt;</b>	Is the command (including arguments) to which you want to assign the specified alias.

For example, to define “?” as an alias for “help,” type:

```
9:ASN-9000:system# alias ? help
Added ?:help
```

The ASN-9000 acknowledges that it has added ? to its list of aliases. To define “hist” as an alias for “history,” issue the following command:

```
10:ASN-9000:sysyem# alias hist history
```

Note that only one level of alias substitution is performed. That is, strings within an alias definition are not checked against the alias list. For instance, in the following example, the “?” alias for help still works even though “help” is defined as an alias for subsystems.

```
11:ASN-9000:system# alias help subsystems
Added help:      subsystems
12:ASN-9000:system# help
atalk atm atm/1483encap atm/clip atm/foreip
atm/lane bridge dec host ip ip/rip ip/
ospf ip/mcast ipx ipx/rip ipx/sap media nvram
snmp system tftp
13:ASN-9000:system# ?
Global commands:
history|hi      show command
history
help listing continues
```

### 5.3.2 Displaying an Alias

To display the definition of an alias, issue the following command:

```
alias [<string>]
```

**<string>**     Is the alias for which you want to display the definition. If you do not specify an alias, all aliases defined for the current session are displayed.

Here is an example of the display produced by this command:

```
14:ASN-9000:system# alias
?      help
help   subsystems
```

### 5.3.3 Saving and Loading an Alias

Aliases apply to the current command-line session only. For example, if you define aliases within a TTY1 session, then open a second TTY session or a TELNET session, the aliases are not available to the new session. Moreover, if you end the TTY1 session without saving the aliases defined during that session, they are lost.

There are several ways to save aliases. The easiest way is to save them to an environment file using the following command:

```
saveenv <file-name>
```

You also can manually add aliases to a file, then type **readenv** <file-name> to read (load) them in each time you log in. Environment files contain other session parameters in addition to aliases.

Alternatively, you can place the aliases in an external file on your management station, then load them into the command-line session each time you log in. For example, if your management station is a PC running Kermit, you can put them in an “aliases” text file on the PC, then load them into the *ForeRunner* ASN-9000 by escaping into the Kermit prompt and typing “transmit aliases”; this command transmits the text file as if you were typing it. On UNIX systems running the tip program, you can use the “~>” escape to send a local text file to the *ForeRunner* ASN-9000.

### 5.3.4 Deleting an Alias

To delete an alias, issue the following command:

```
unalias <name>
```

<name>      Is the alias you want to delete.

Here is an example of the use of this command. In this example, the “?” is a string that was defined as an alias by an **alias** command:

```
15:ASN-9000:system# unalias ?
```

## 5.4 Using Timed Commands

---

When you use a command-line session to monitor network behavior, you might want to execute certain commands regularly and repeatedly. For example, you might want to display a bridge or route cache at regular intervals to observe frequently requested bridge or route destinations for certain segments.

You can define and activate a timed command to automatically issue any command string at a regular interval. A *timed command* is similar to an alias, but is automatically issued by the *ForeRunner* ASN-9000 at an interval you specify. You can define *ForeRunner* ASN-9000 commands and even aliases as timed commands. Each user session can have up to eight timed commands.

The following lists the commands used to define, display, activate, and delete timed commands.

**Table 5.1 - Timed Commands**

Use Command...	To...
<code>timedcmd   tc</code>	Display all timed commands.
<code>timedcmd   tc add &lt;id&gt; &lt;interval(secs)&gt; &lt;cmd-and-args&gt;</code>	Add a timed command.
<code>timedcmd   tc enable &lt;id&gt;</code>	Start the timer for a timed command.
<code>timedcmd   tc disable &lt;id&gt;</code>	Stop the timer for a timed command.
<code>timedcmd   tc del &lt;id&gt;</code>	Delete a timed command.

Each user session can have a maximum of eight timed commands. As with command aliases, timed commands are local to the current command-line session. That is, they are not remembered across logins or resets of the *ForeRunner* ASN-9000 unless you save them to an environment file. Also, timed commands are automatically canceled when the command-line session ends. Timed commands can be stored in an environment file by issuing the `saveenv <filename>` command.

### 5.4.1 Defining a Timed Command

Use the `timedcmd add` command to define a timed command. Here is the syntax for this command:

```
timedcmd | tc add <id> <interval(secs)> <cmd-and-args>
```

- <id>** Specifies the name of the timed command. When you activate the timed command, you use this name. You can specify an alphanumeric string up to 15 characters in length.
- <interval (secs)>** Specifies, in seconds, the interval at which the timed command is reissued. You can specify a minimum of 1 second.
- <cmd-and-args>** Specifies the command string that is issued each time the interval specified by (secs) expires. You can specify a *ForeRunner* ASN-9000 command, including its arguments, or an alias.

**NOTE**

You must include the subsystem name in each timed command. For example, if you create a timed command that issues the **interface** command from within the **atalk** (AppleTalk) subsystem, specify the command as **atalk interface**.

Here is an example of how to define a timed command. In this example, a timed command named “**bcache**” is defined to automatically display the bridge cache every ten seconds. A command such as this is useful for quickly observing bridge activity.

```
35:ASN-9000:system# timedcmd add bcache 10 bridge display-cache 1-6
Added bcache: 10 secs, bridge display-cache 1-6 (timer not running)
```

## 5.4.2 Starting a Timed Command

To use a timed command, you must start the timer for the command. When you start the timer, the command is issued when the specified timer value expires. The *ForeRunner* ASN-9000 continues to issue the timed command each time the interval expires until you issue the **timedcmd disable** command, or until you log out. Note that if you issue the **system saveenv <file-name>** command while the timed command is running, the **timedcmd enable <id>** command is added to the environment file. Consequently, the timed command is started again the next time you read the environment file specified by **<file-name>**.

Use the **timedcmd enable** command to start the timer for a timed command. Here is the syntax for this command:

```
timedcmd|tc enable <id>
```

**<id>** Specifies the name of the timed command for which you are starting the timer. Make sure you specify the name of the timed command itself, rather than the command string associated with the timed command.

## 5.4.3 Stopping a Timed Command

Use the **timedcmd disable** command to stop a command timer. When you stop a command timer, the *ForeRunner* ASN-9000 software stops issuing the timed command. Here is the syntax for the **timedcmd disable** command.

```
timedcmd|tc disable <id>
```

**<id>** Specifies the name of the timed command. Make sure you specify the name of the timed command itself, rather than the command string associated with the timed command.

You can stop a timed command by ending the command-line session from which the timed command was started.

You can restart the timed command by issuing the following command: `timedcmd enable <id>`.

If you save an environment file while the timed command is running, or manually add the timed command (and the `timedcmd enable <id>` command) to the environment file, the timed command starts again when you read (load) the environment file.

## 5.4.4 Deleting a Timed Command

Use the `timedcmd del` command to delete a timed command. Here is the syntax for this command:

```
timedcmd|tc del <id>
```

**<id>** Specifies the name of the timed command you want to delete.

Here is an example of the use of this command:

```
37:ASN-9000:system# timedcmd del bcache
bcache: deleted
```

## 5.5 Using Environment Files

At any time during a command-line session, you can save or read (load) an *environment file*, an ASCII file that contains *ForeRunner* ASN-9000 commands defining the following parameters:

- Scroll (stty) parameters (maximum number of rows and “more” enable or disable; defined using the `stty` command).
- Command aliases (created using the `alias` command).
- Timed commands (created using the `timedcmd` command).

Normally, when you end a session by logging out or rebooting, any changes made to these environment settings are lost. However, if you save an environment file before logging out or rebooting, environment settings are placed into an environment file.

## More Global Commands

If you read (load) that file during a command-line session, the commands in the file recreate the scroll parameters, command aliases, and timed commands that were active when you saved the file.

Here is an example of an environment file:

```
#
# stty
system stty rows 24
system stty more enl
#
# aliases
#
system alias aarp      atalk at *.1
system alias br        bridge s all pi,po
system alias bru       bridge s all pu,cu
system alias bsc       bridge sc
system alias bt        bridge bt
system alias btc       bridge btc
system alias ethan     ip ping 181.17.45.17
system alias sascha    ip ping 191.1.45.3
system alias my        ip ping 131.24.45.2
system alias boys     ip ping 131.24.45.3
system alias sat       atalk s ddp
system alias si        ip s ip
system alias sx        ipx s ipx
#
# timed commands
#
system timedcmd add bcache 10 bridge display-cache 1-6
```

Notice that the file contains three sections: **stty**, **aliases**, and **timed** commands, shown in bold type in the example. The sections are labeled by comment lines, which begin with #. The **stty** section contains commands to set the “more” feature and specify the maximum number of rows to be displayed when the “more” feature is enabled. The **alias** section contains **alias** commands that define various aliases. In this example, aliases are created to display tables and packet statistics, and to ping IP addresses. The **timed** commands section contains a command that defines the timed command **bcache**.

You can save or load multiple environment files during a command-line session; the effects are cumulative. Note, however, that earlier settings can be overwritten by later settings. For example, if you read a file that contains the **stty +more** command (to enable the “more” feature), then read another file that contains the **stty -more** command, the net effect will be that “more” is disabled. If you experience unexpected results when using multiple environment files, examine the files to ensure that you are not inadvertently overwriting some parameters.

You can cause an environment file to be automatically loaded at the beginning of a command-line session by saving the file as one of the following:

**root.env** If the ASN-9000 is booted from the floppy drive or the Flash Memory Module, and either the Lock Switch is off or you log in under root capability, this file is loaded if present.

**monitor.env** If the ASN-9000 is booted from the floppy drive or the Flash Memory Module, the Lock Switch is on, and you log in under monitor capability, this file is loaded if present.



If you choose to manually edit an environment file, do not place into the file any commands associated with using the floppy drive or the Flash Memory Module. This includes commands such as **readenv** and **readcfg**, which read files.

Also, if the environment file contains the **timedcmd enable** command and the timed command it starts has been defined in the environment file or earlier in the user session, the timed command is started when the environment file is read.

The following sections describe how to save the environment settings and how to activate the settings during a user session.

## 5.5.1 Saving an Environment File

To save an environment file, issue the following command:

```
saveenv | svenv <file-name>
```

**<file-name>** Specifies the name of a floppy file to contain the environment parameters. You can specify up to eight alphanumeric characters plus a three-character extension. Use a period to separate the name from the extension. It is recommended that you always use the extension **env** (for example: **Lab1.env**) to distinguish environment files from other files.



**NOTE**

If you are in monitor mode, you cannot save to the file name `root.env`. You are in monitor mode if you logged on to the user session as “monitor.”

Also, this command assumes that you want to write the environment file to the default local storage device. You can explicitly specify the storage device by prefacing the file name with **fd:** (floppy drive) or **fm:** (Flash Memory Module).

Here is an example of how to save a default environment file. In this example, the command-line session is in monitor capability. Correspondingly, the current environment settings are saved into a file called `monitor.env`:

```
32:ASN-9000:system> saveenv monitor.env
monitor.env: Environment saved
```

In the following example, the scroll parameters are set, then some aliases and a timed command are defined. These environment settings are then saved into an environment file named `Lab1.env`. At any time during a command-line session, this file can be read (loaded) to activate the environment settings it contains.

*All* current environment settings (aliases, timed commands, and scroll parameters) are saved in the environment file. These commands create the environment file shown in the following example. As you can see, environment files can save you a lot of typing:

```
33:ASN-9000:system# stty rows 24
34:ASN-9000:system# stty -more
35:ASN-9000:system# alias aarp    atalk at *.1
36:ASN-9000:system# alias br     bridge s all pi,po
37:ASN-9000:system# alias bru    bridge s all pu,cu
38:ASN-9000:system# alias bsc    bridge sc
39:ASN-9000:system# alias bt     bridge bt
40:ASN-9000:system# alias btc    bridge btc
41:ASN-9000:system# alias clr    media rdcfg clearASN-9000
42:ASN-9000:system# alias ethan  ip ping 181.17.45.17
43:ASN-9000:system# alias sascha ip ping 191.1.45.3
44:ASN-9000:system# alias bill   ip ping 131.24.45.2
45:ASN-9000:system# alias ginger ip ping 131.24.45.2
46:ASN-9000:system# alias sat    atalk s ddp
47:ASN-9000:system# alias si     ip s ip
48:ASN-9000:system# alias sx     ipx s ipx
49:ASN-9000:system# timedcmd add bcache 10 bridge display-cache 1-6
50:ASN-9000:system# saveenv monitor.env
Lab1.env: Environment saved
```

## 5.5.2 Reading (Loading) an Environment File

You can load an environment file and thereby activate the environment settings saved in that file using the following command:

```
readenv|rdenv <file-name>
```

**<file-name>** Specifies the name of the file that contains the environment settings. If the file is found, the commands in the file are executed; otherwise, an error message is displayed.

### NOTE

This command assumes that the environment file is located on the default local storage device. You can explicitly specify the storage device by prefacing the file name with **fd:** (floppy drive) or **fm:** (Flash Memory Module).

## 5.5.3 Editing an Environment File

As an alternative to setting environment parameters, then using **saveenv** to create or edit an environment file, you can edit the file directly using a text or ASCII editor.

### NOTE

Do not place any commands associated with using the floppy drive or Flash Memory Module in the file. This restriction includes commands such as **readenv** and **readcfg**, which read files.



# CHAPTER 6

## System Subsystem Commands

This chapter describes the commands in the system subsystem. The system subsystem commands control various system-level settings on the *ForeRunner* ASN-9000. The following tasks can be performed using these commands:

- Set or display the baud rate configured on TTY1 and TTY2.
- Display the date and time the ASN-9000 was last loaded and the boot order.
- Remove and install modules while the ASN-9000 is operating.
- Display the current configuration
- Set or display the system date and time.
- Enable or disable automatic segment state detection.
- Read the MAC-layer hardware address.
- Display prom information of installed modules.
- Set or change management level passwords.
- Read the default or an alternate configuration file.
- Reboot the ASN-9000.
- Save the default or an alternate configuration file.
- Assign a location for the ASN-9000.
- Assign a name to the ASN-9000.
- Display the temperature of one or all modules installed.
- Enable or disable the TTY2 port.
- Display the elapsed time since the last reboot.
- Display the software version string and idprom info of the card in the specified slot or all the cards in ASN-9000.

### 6.1 Accessing the System Subsystem

---

The `system` subsystem is the default subsystem entered when the ASN-9000 completes the boot process. To access the `system` subsystem from any other subsystem, enter **system** from the current runtime prompt.

## 6.2 System Commands

---

There are several commands in the system subsystem that give information on the state and configuration of the ASN-9000. These commands allow configuration and environmental information about the ASN-9000 to be displayed.

### 6.2.1 Changing the TTY Port Baud Rate

The baud rate (data transmission rate) for the RS-232 ports (tty1 and tty2) can be changed using the **baud** command. However, before setting the baud rate associated with tty2, the tty2 port must be enabled using the **tty2** command (see Section 6.2.15). The syntax for this command is:

```
baud set tty1|tty2 1200|2300|4800|9600|19200
baud [show]
```

<b>set</b>	Sets the specified baud rate for the specified port.
<b>tty1 tty2</b>	Specifies the port to be set.
<b>1200 2300 4800 9600 19200</b>	Select one of the listed baud rates to apply to the specified port.

The newly specified rate is stored in NVRAM and takes effect immediately. It is retained across logins and power cycles. An example of the use of this command is:

```
43:ASN-9000:system# baud
TTY          Baud Rate
1            9600
2            19200
44:ASN-9000:system# tty2 enable
45:ASN-9000:system# baud set tty2 9600
Changed tty2 baud rate to 9600; written to nvram
46:ASN-9000:system# tty2 disable
tty2 is now closed
47:ASN-9000:system#
```



If the Lock Switch is unlocked when booting the ASN-9000, the TTY ports use the default baud rates (9600 for TTY1 and 1200 for TTY2), regardless of the baud rates stored in NVRAM.

## 6.2.2 Removing and Replacing Interface Modules

Whenever it is necessary to remove a NIM in the ASN-9000 chassis, the **card-swap disable** command must be issued so that the ASN-9000 configuration manager can deactivate traffic to ports/segments on that card. Conversely, when replacing a NIM, of the same type in the same slot, the **card-swap enable** command reactivates traffic to the ports/segments on the card. These operations can be accomplished while the ASN-9000 is operating. Entering the **card-swap** command with no parameters displays a list of the modules currently installed in each card slot.

### CAUTION



The **card-swap** command is only to be used with interface modules (NIMs). Do not use this command to change out the Packet Engine. The ASN-9000 must be powered down to change out the Packet Engine. Removing the Packet Engine, while the ASN-9000 is operating can result in damage to the Packet Engine.

### NOTE



NIMs can only be swapped when the chassis contains at least one redundant power module. (See the *ForeRunner ASN-9000 Hardware Reference Manual* for information about power redundancy.)

### NOTE



The NIM being installed must be of the same type as was removed. If a NIM is being replaced with a different type, the proper procedure is to power down the ASN-9000, remove the card, insert the new card and then power on the ASN-9000. This sequence loads the ID PROM information of the cards currently installed into the ASN-9000 configuration manager.

The syntax for this command is:

```
card-swap|cs enable|disable <slot> card-swap|cs [show]
```

**enable|disable**

Entering **disable** states that the card is being removed from the ASN-9000. **Enable** states that the card is being restored to the ASN-9000.

**<slot>** Specifies which card slot is being enabled/disabled.

Refer to the *ForeRunner ASN-9000 Hardware Reference Manual* for detailed procedures on removing and installing NIMs in the ASN-9000 chassis.



There may be a slight delay when issuing the **card-swap enable** command depending on the type of card. This is to allow the ASN-9000 to reload a required image to the newly installed card.

When this procedure is complete, the segments on the replaced card are reactivated and traffic starts passing traffic normally.

In the following examples, the NIM in slot 1 (the bottom slot in the chassis) is removed using the **card-swap disable** command. Notice that when a NIM slot is empty, the **config** display does not show a card in that particular slot.

```
67:ASN-9000:system# config
Accelerator board is present. Accelerator IOP is being used.
Installed DRAM Size: 24 MB
tty1: 9600 baud
tty2: 9600 baud
PE: slot 5
PM1: present and good
PM2: not present
PM3: not present
PM4: not present
```

```
03/33 MM/MM
02/17 UTP      UTP      UTP      UTP      UTP      UTP
      UTP      UTP      UTP      UTP      UTP      UTP
      UTP      UTP      UTP      UTP
01/01 OC3-MF  OC3-MF  OC3-MF  OC3-MF  OC3-MF  OC3-MF
      OC3-MF  OC3-MF  OC3-MF  OC3-MF  OC3-MF  OC3-MF
      OC3-MF  OC3-MF  OC3-MF  OC3-MF  ----  ----
      ----  ----  ----  ----  ----  ----
      ----  ----  ----  ----  ----  ----
      ----  ----
```

```
68:ASN-9000:system# card-swap disable 1
Card 1 removed.
69:ASN-9000:system# config
Accelerator board is present. Accelerator IOP is being used.
Installed DRAM Size: 24 MB
tty1: 9600 baud
tty2: 9600 baud
PE: slot 5
PM1: present and good
PM2: not present
```

```

PM3:    not present
PM4:    not present
03/33 MM/MM
02/17 UTP    UTP    UTP    UTP    UTP    UTP
      UTP    UTP    UTP    UTP    UTP    UTP
      UTP    UTP    UTP    UTP
70:ASN-9000:system# card-swap enable 1
GINIM BOOT: slot 1, image "FM:7atm"
Card 1 inserted.
71:ASN-9000:system#

```

The following example shows the response if the **card-swap disable <slot>** command is issued against an empty slot:

```

71:ASN-9000:system# card-swap disable 4
Card 4 was not present during boot up time.
72:ASN-9000:system#

```

In the following example, a NIM has been inserted in the specified slot, yet the thumbscrews on the NIM have not been tightened. As a result, the left ejector handle is not pressing on the activation switch. The NIM fails to function until the activation switch is pressed by the ejector handle. The activation switch is located behind the left ejector handle. The ejector handle must be depressed fully, and the thumbscrews tightened before the activation switch is properly enabled.

```

74:ASN-9000:system# card-swap disable 3
Card 3 removed.
75:ASN-9000:system# card-swap enable 3
Card 3 has not been fully latched. Please ensure that the thumb crews on
the card are fully tightened and reissue this command.
76:ASN-9000:system#

```

## 6.2.3 Displaying the ASN-9000

To display the current configuration of the ASN-9000, issue the **config** command. This command displays information on the physical configuration of the ASN-9000. The following is an example of the **config** command. In this example, information is displayed for a ASN-9000 5-slot chassis. Note that because no segments were specifically allocated for the PowerCell NIM, which starts with segment 02/07, only six virtual interfaces were configured. The PowerCell 700 supports up to 32 virtual interfaces.

```

3:ASN-9000:system# config
Accelerator board is present. Accelerator IOP is being used.
Installed DRAM Size: 24 MB
tty1:  9600 baud
tty2:  19200 baud
PE:    slot 5
PM1:   present and good
PM2:   not present
PM3:   not present

```

## System Subsystem Commands

```
PM4:    not present
03/33 MM/MM
02/17 UTP      UTP      UTP      UTP      UTP      UTP
      UTP      UTP      UTP      UTP      UTP      UTP
      UTP      UTP      UTP      UTP
01/01 OC3-MF   OC3-MF   OC3-MF   OC3-MF   OC3-MF   OC3-MF
      OC3-MF   OC3-MF   OC3-MF   OC3-MF   OC3-MF   OC3-MF
      OC3-MF   OC3-MF   OC3-MF   OC3-MF   ----
      ----
      ----
      ----
      ----
4:ASN-9000:system#
```

The **config** command displays the following information.

- Whether a Packet Accelerator is present on the Packet Engine and if the accelerator input/output processor (IOP) is in use or not.
- The amount of dynamic random access memory (DRAM) on the Packet Engine.
- The current baud rates for RS232 tty1 and tty2 ports.
- The slot occupied by the Packet Engine, indicated by PE:. In this example, the Packet Engine is in slot 5, the top slot in a 5-slot chassis.
- The presence and status of power module, indicated by PM1:, PM2:, PM3:, and PM4:.
- The slot number and starting segment number for each slot, and the media type in use in each segment position. The row beginning 01/01: displays the configuration of the NIM in slot 1, beginning with segment 1. The row beginning 02/17: shows the configuration in slot 2, beginning with segment 17, and so on. Remember that segments are numbered from left to right, bottom to top. The number of segments in each NIM slot depends on how many are allocated in NVRAM. If slots are not explicitly allocated in NVRAM, the ASN-9000 defaults to the maximum allowable segments per NIM. Empty NIM slots are not displayed.

### 6.2.4 Setting and Displaying the System Time and Date

The **date** command is used to display or set the system date and time. The syntax for this command is:

```
date set [YYMMDD]hhmm[.ss] date [show]
```

<b>set</b>	Sets the specified date and/or time.
<b>[YYMMDD]hhmm&gt;</b>	Specifies the year (YY), month (MM), day (DD), hour (hh), and minute (mm). To set the time, but not the date, specify <hhmm>[.<ss>]. (The software reads this argument from right to left, so any additional

arguments can be specified with `<hhmm>`. For example, specifying `<DDhhmm>`, also specifies the day. Note that the arguments must be specified in the order shown. For example, `<YYhhmm>` or `<DDMMYYhhmm>` cannot be entered.)

**[.ss]** Optionally specifies the seconds. If this argument is used, make sure to use the period (.) in front of the seconds. If the number of seconds is not specified, the value is set to 00.

If either argument is not used, the current system date and time are displayed.

Here are some examples of the use of this command. In the first example, the current system time and date are displayed. In the second example, the time and date are changed. The new time and date are then displayed.

```
29:ASN-9000:system# date
Wed Aug 27 11:27:15 1997
30:ASN-9000:system# date set 9708271130.10
date set to: Wed Aug 27 11:30:10 1997
31:ASN-9000:system# date
Wed Aug 27 11:30:15 1997
32:ASN-9000:system#
```

## 6.2.5 Setting the Data Carrier Detect Parameter

The Data-Carrier Detect parameter can be displayed or changed using the `dcd-detection` command. Here are some examples of the command. The syntax for this command is:

```
dcd-detection|dcd enable|disable dcd-detection|dcd [show]
```

**enable|disable** Specifies whether to enable or disable data-carrier detection

Note that in this example, the short form of the `dcd-detection` command is used.

```
48:ASN-9000:system# dcd
dcd-detection is currently disabled.
49:ASN-9000:system# dcd enable
dcd-detection enabled
50:ASN-9000:system#
```

## 6.2.6 Displaying the ASN-9000 MAC Address

To display the Mac address of the ASN-9000, issue the **ethaddr** command. Here is the syntax for this command:

```
ethaddr|ea [show]
```

Here is an example of the display produced by this command. In this example, the short form of the command is used.

```
9:ASN-9000:system# ea
Ethernet address: 00-00-ef-03-9a-b0
10:ASN-9000:system#
```

## 6.2.7 Displaying ID and Power Information

The Packet Engine and all types of NIMs contain a special PROM called the ID PROM. The ID PROM contains identification information and power requirements for the module. This information can be displayed using the **idprom** command. Here is the syntax for this command:

```
idprom|idp [show] <slot number>|all
```

**<slot number>** Specifies the NIM slot containing the module.

Here is an example of the display produced by this command. In this example, information is displayed about the NIM in slot 2.

```
54:ASN-9000:system# idprom 2

Card Type: UTP 16x1 Interface Module
Serial #: 632027371
Model: 7202-00
Revision: G
Issue: 1
Deviation: <not set>

Power Requirements:
5000 mA at 5V
55:ASN-9000:system#
```

The ID PROM display shows the following information:

<b>Card Type:</b>	Specifies the card currently installed in the slot specified.
<b>Serial #:</b>	Displays the serial number of the specified card.
<b>Model:</b>	Displays the model number of the specified card.

<b>Revision:</b>	Displays the revision level of the card.
<b>Issue:</b>	Displays the card issue number.
<b>Deviation:</b>	If applicable, displays the factory-assigned deviation number. Only some modules have deviation numbers.
<b>Power Requirements:</b>	Displays the maximum amperage (milliamps) required by the module at +12-volts, +5-volts, or +3.3-volts, as applicable.

If the **idprom** command is issued against an empty card slot, the following message is displayed:

```
55:ASN-9000:system# idprom 4
unable to read IDPROM information from slot 4
56:Asn-9000:system#
```

## 6.2.8 Changing the Password

The **passwd** command is used to change the system password associated with “root” or “monitor” logins. The syntax for this command is:

```
passwd [root|monitor]
```

**root|monitor** Indicates the management capability for which the password is being changed.

To change a password:

1. Issue the **passwd** command, specifying the appropriate management level (root or monitor) capability. A prompt is displayed to enter the new password.
2. Enter the new password to be assigned to this management level. If no password is to be set, press Enter.
3. A prompt is then presented to Re-enter the password (Re-enter new password:) previously entered.
4. Re-enter the password that was entered at the New password: prompt, press Enter.



This prompt is not displayed if the Lock Switch is in the unlocked position (U) or the Lock Switch jumper is set to Unlock. Instead, the New password: prompt is displayed.

5. The message “Password changed” is displayed to confirm that the password was changed.

The following example shows how a password for root management capability is changed.

```
52:ASN-9000:system# passwd root
New password:*****
Re-enter new password:*****
Password changed
53:ASN-9000:system#
```

For security reasons, the input shown above with asterisks does not appear when entered in response to the prompts. Remember that passwords are not required if the Lock Switch is in the unlocked (U) position. If the password is forgotten, turn the Lock Switch off, log in and enter a new password, then turn the Lock Switch on again.

### 6.2.9 Reading a Configuration File

When loading the software, the system looks for a configuration file on the default-device:

- If the Flash Memory module was used to load the software, the system looks for a file named `cfg`. If present, this file is automatically loaded and its configuration information is used to configure the ASN-9000.
- If loaded from a BOOTP/TFTP server, the system looks for the configuration file specified in the `bootdef` (boot definition) file on the server.

Even if the system finds and loads a configuration file when the software is booted, additional configuration files can be loaded during a ASN-9000 session using the `readcfg` command.



The new configuration information does not undo the configuration information contained in the `cfg` file. Instead, the new configuration is added to the current configuration, until the ASN-9000 is powered down or rebooted. The additional, or different, configuration information can be saved with the current configuration information by issuing the `savecfg` command.

### 6.2.9.1 Loading a Configuration from Flash Memory

To load a configuration file located in Flash Memory (or a saved on a terminal connected to a TTY port), issue the following command:

```
readcfg|rdcfg [-v] <file or device name>
```

**-v** Directs that each line of the configuration file be displayed to the console during execution.

**<file or device name>** Specifies a file name is to be read from the device specified. If no device is specified the default-device is assumed. For auto-configuration on boot up, use the file name 'cfg'.

#### NOTE

The last line in any configuration file must be the string 'endcfg' or 'ecfg'.

### 6.2.9.2 Loading a Configuration From a TFTP Server

To load a configuration file located on the default TFTP server, issue the following command:

```
readcfg|rdcfg [-v] [-h <host>] <remote-file>
```

**-v**

**-h <host>** Specifies the IP address of the TFTP server. Unless a default TFTP server was specified using the tftp set command, this argument must be included.

**<remote-file>** Specifies the configuration file name. Specify a name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called fore and this directory is specified as the TFTP home directory, do not specify fore as part of the file name.

**NOTE**

The tftp version of the **readcfg** command only works if the system was last boot from a tftp boot server.

## 6.2.10 Rebooting the ASN-9000

The ASN-9000 can be loaded from a command-line session by entering the **reboot** command. The **reboot** command performs a cold restart of the ASN-9000. During a cold restart, the Packet Engine conducts a power-on self-test to check its various hardware components. Following successful completion of the power-on self-tests, the ASN-9000 software is loaded.

## 6.2.11 Saving the Configuration

Configuration changes can be saved to the configuration file (**cfg**) in Flash Memory (or on the TFTP server), in which case they are automatically applied each time the software is loaded. Alternatively, changes can be saved to a different filename or to a device attached to TTY1 or TTY2. If configuration changes are saved to a file other than **cfg**, the file must be loaded after the software is loaded to apply the configuration settings to the ASN-9000.

### 6.2.11.1 Saving Configuration Files to Flash Memory

To save a configuration file, issue the following command:

```
savecfg|svcfg <file or device name>
```

**<file or device-name>** A filename must be specified. If a device is not specified the configuration file is saved to the default-device.

In the following example, the current configuration is saved to a file named **Lab1.cfg** on the default-device.

```
98:ASN-9000:system# savecfg Lab1.cfg
99:ASN-9000:system#
```

### 6.2.11.2 Saving the Configuration to a TFTP Server

To save a configuration file to a TFTP server, issue the following command:

```
savecfg|svcfg [-h <host>] <remote-file>
```

**-h <host>** Specifies the IP address of the TFTP server, if different than the address specified with the **tftp set server** command. Unless a TFTP server was specified using the **tftp set server** command,

include this argument. For information on the **tftp** commands, refer to the *ForeRunner ASN-9000 Software Reference Manual*.

**<remote-file>**

Specifies the name of the configuration file to be saved. Specify a name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called *fore* and this directory is specified as the TFTP home directory, do not specify *fore* as part of the file name.

**NOTE**

The tftp version of the **readcfg** command only works if the system was last boot from a tftp boot server. Refer to the *ForeRunner ASN-9000 Software Reference Manual* for instructions on setting up a TFTP boot server.

**NOTE**

Some TFTP servers require that the remote file name exist on the server before writing to that file name. If the server requires that the filename already exist, create a short file (named the same as the configuration file) on the server, then specify that file name for *<remote-file>*.

**NOTE**

On some TFTP servers, including servers running Sun/OS 4.x, files overwritten on the server are not properly truncated. When overwriting an existing file on the TFTP server, if the older version of the file is longer than the new file, the older version is not truncated properly by the server. As a result, the new version of the file contains part of the older version of the file.

An example of the use of this command is.

```
99:ASN-9000:system# tftp svcfg 147.128.128.7 Lab1.cfg
Configuration saved to Lab1.cfg
```

## 6.2.12 Setting and Displaying the System Location

The system location can be changed using the **syslocn** command. The syntax for this command is:

```
syslocn set <location> syslocn [show] <location>
```

<b>set</b>	Sets the location specified.
<b>&lt;location&gt;</b>	Specifies the location of the ASN-9000. Any alphanumeric string up to 24 characters in length can be specified. The location name cannot contain blanks.  If a location is not specified, the current location name is displayed.

The following example shows how to display the current system location and to change the location variable. The new system location is defined as “Pittsburgh.” (Note that the system name is now “ASN-9000,” as changed by the **sysname** command in the example above.)

```
26:ASN-9000:system# syslocn  
Current system location is: Undefined  
  
27:ASN-9000:system# syslocn set Pittsburgh  
System location set to:  
Pittsburgh  
28:ASN-9000:system#
```

## 6.2.13 Setting and Displaying the System Name

The system name is shown in the command prompt. The default system name is ASN-9000. The system name can be changed using the **sysname** command. The syntax for this command is:

```
sysname set <location> sysname [show] <location>
```

<b>set</b>	Sets the name specified.
<b>&lt;location&gt;</b>	Specifies the name assigned to this ASN-9000. Any alphanumeric string up to 24 characters in length can be specified. The name cannot contain blanks.  If a name is not specified, the current name is displayed.

The following example shows how to display the current system name and to change the name variable. The new system name is defined as “ASN-9000.”

```

22:ASN-9000:system# sysname
Current system name is: ASN-9000
23:ASN-9000:system# sysname set ASN-9000
System name set to 'ASN-9000'.
24:ASN-9000:system#

```

## 6.2.14 Displaying the Temperature of the ASN-9000

The Packet Engine and NIMs contain a temperature sensor that reads the temperature of the module with an accuracy of plus or minus 0.5° C. The current temperature of a module can be displayed by issuing the following command:

```
temperature|temp [show] <slot number>|all
```

**<slot number>** Specifies the slot that contains the module for which to display the temperature.

**all** Displays the temperature for all installed modules.

In the example that follows, the first command displays the temperature for all modules. The second, displays the temperature for the module in slot 2.

```

11:ASN-9000:system# temp all
slot 5, temp 41.5 degrees C
slot 3, temp 35 degrees C
slot 2, temp 32 degrees C
slot 1, temp 0 degrees C
12:ASN-9000:system# temp 2
slot 2, temp 32 degrees C
13:ASN-9000:system#

```

Note that the ASN-9000 is designed to operate over a range of external ambient temperatures. An additional temperature rise inside the chassis is accounted for in the design of the product.

Some older revisions of the Packet Engine and NIMs do not contain an ID PROM. If the **temperature** command is issued against a module that does not contain an ID PROM, or against a slot that does not contain a NIM, the system displays the following message:

```

13:ASN-9000:system# temp 4
slot 4, temp not available
14:ASN-9000:system#

```

## 6.2.15 Enabling/Disabling TTY2

Before setting or changing the baud rate associated with tty2 using the **baud** command (see Section 6.2.1), the tty2 port must be enabled using the **tty2** command. The syntax for this command is:

**tty2 enable|disable**

**enable|disable**      Enables or disables the tty2 port.

An example of the use of this command is:

```
43:ASN-9000:system# baud
TTY          Baud Rate
1            9600
2            19200
44:ASN-9000:system# tty2 enable
45:ASN-9000:system# baud set tty2 9600
Changed tty2 baud rate to 9600; written to nvram
46:ASN-9000:system# tty2 disable
tty2 is now closed
47:ASN-9000:system#
```



If the Lock Switch is unlocked when booting the ASN-9000, the TTY ports use the default baud rates (9600 for TTY1 and 1200 for TTY2), regardless of the baud rates stored in NVRAM.



Before the baud rate can be set for TTY2, a session must be opened on the port. To open a TTY2 session, issue the **tty2 enable** command.

## 6.2.16 Displaying the System Uptime

The **system uptime** command displays how long much time has elapsed since the last time the ASN-9000 was loaded. There are no parameters on for the **uptime** command. Here is an example of the **uptime** command.

```
15:ASN-9000:system# uptime
Elapsed time since last reboot: 3 hours, 9 minutes, 9 seconds
16:ASN-9000:system#
```

## 6.2.17 Displaying the ASN-9000 Software Version

The version command displays the version level of software currently running on the NIMs and Packet Engine installed in the ASN-9000. The syntax of this command is:

<b>version ver [show] [&lt;slot-number&gt; all]</b>	
<b>&lt;slot-number&gt;</b>	Lists the version of software on the card in the slot specified.
<b>all</b>	Lists the version of software on the Packet Engine and all Intelligent NIMs installed in the ASN-9000. If the all parameter is not used, software version information is displayed for the Packet Engine only.

A typical display of the use of this command is shown below.

```
61:ASN-9000:system# ver

Card Type: Packet Engine - 40MHz
Serial #: 633020265
Model: 7101-01
Revision: C
Issue: 2
Deviation: <not set>

ASN-9000 Version: PH 7pe FT4.0.0 (@5567) (s#5) 1997.08.23 17:14
PROM Version: 7pep-2.5.5 (sl.85) 1996.06.21 11:14
62:ASN-9000:system#
```

<b>Card Type:</b>	Specifies the card currently installed in the slot specified.
<b>Serial #:</b>	Displays the serial number of the specified card.
<b>Model:</b>	Displays the model number of the specified card.
<b>Revision:</b>	Displays the revision level of the card.
<b>Issue:</b>	Displays the card issue number.
<b>Deviation:</b>	If applicable, displays the factory-assigned deviation number. Only some modules have deviation numbers.
<b>ASN-9000 Version:</b>	Displays the ASN-9000 type, software version installed, with the software build number and date and time of build.

**PROM Version:** Displays the PROM version information which includes the version number with build and date and time of build.

## 6.2.18 Displaying Boot Information on the ASN-9000

The ASN-9000 software can be loaded from floppy diskette, Flash Memory Module. After the software is loaded, the following information is logged in memory as the bootlog. This information can be retrieved by executing the **bootinfo** command. The bootlog contains:

- The date and time the system was started.
- The date, time and nvram bootorder (see *Chapter 7, NVRAM Subsystem Commands* on setting the nvram boot order).
- The boot device used to boot. The value can be f (floppy diskette), m (Flash Memory Module), or n (network). This value shows the boot source actually used, which may differ from the boot order specified in NVRAM.

The **bootinfo** command displays the contents of the bootlog. An example of the **bootinfo** command is.

```
7:ASN-9000:system# bootinfo
Wed Aug 27 08:00:57 1997 start
Wed Aug 27 08:01:02 1997 nvram boot order: m
boot device: m
8:ASN-9000:system#
```

## 6.3 System Configuration Commands

---

The system subsystem contains commands that allow the changing of configuration parameters. The following sections provide more information on changing configuration information.

### 6.3.1 Rebooting Without Loading the Default Configuration File

When rebooting the software, the system looks in the bootdef file for the name of a configuration file. If the software needs to be loaded without loading the configuration file, use the global command **rename** command to temporarily rename the default configuration file before loading. For example, if the boot definition file calls the configuration file **cfg**, rename **cfg** to **temp.cfg**. The syntax for this command is:

```
mv|rename <file1> <file2>
```

**<file 1>** Specifies the current filename.

**<file 2>** Specifies the new filename.

After loading the software, restore the configuration file to its original name using the same command. When loading a second time, the ASN-9000 processes the `cfg` file and boots with the default configuration.

## 6.3.2 Editing a Configuration File

Although changes to a configuration file can be automatically written using the `savecfg` command, the configuration file can be edited manually using a text or ASCII editor. To move the file to a machine that contains an ASCII editor:

- If the editor is on a file server, use the **tftp put** command to write the file to the server, edit the file, then use the **tftp get** command to download the edited file back into the Flash Memory Module. (See *ForeRunner ASN-9000 Software Reference Manual*, for information on **tftp** commands.)
- If the editor is on a PC or Macintosh, use the ZMODEM **zw** command to write the file to the PC or Macintosh, edit the file, then use the **zr** command to transfer the edited file to the Flash Memory Module. These commands are available from the boot PROM prompt, <PROM-7pe>.

## 6.3.3 Capturing Configuration Information

The system software diskettes contain a file that enables configuration information about the ASN-9000 system to be captured. The file is called `dispcfg`. When reading the file (using the **readcfg dispcfg** command), the commands in the file display configuration information on the management terminal.

If problems are experienced, FORE Systems TAC (Technical Assistance Center) might request that the `dispcfg` file be read so that the information captured by the file may be helpful in resolving the problems. To read the `dispcfg` file, issue the following command:

```
system readcfg dispcfg
```



# CHAPTER 7

## NVRAM Subsystem Commands

This chapter explains the NVRAM subsystem and the commands used to make changes in the NVRAM subsystem and booting parameters. The commands in the NVRAM subsystem affect the following ASN-9000 parameters:

- Boot source order
- ASN-9000 IP address
- ASN-9000 local subnet mask
- Gateway address (when a gateway (router) separates the ASN-9000 from a BOOTP server)
- A TFTP server address
- Post-system crash behavior
- Segment allocation on NIM slots
- RIPv2 encryption key

### 7.1 NVRAM Configuration Commands

---

The NVRAM commands can be used to show, set, change or clear the parameters listed above. This section explains the commands used to set NVRAM variables using commands in the `nvr` subsystem. The commands in the `nvr` subsystem can be used to set the values used by the ASN-9000 for:

- Setting the boot order.
- Setting IP addresses to identify this ASN-9000, default file server and gateway.
- Behavior following an unexpected system crash.
- The number of segments allocated to each NIM slot.

## 7.1.1 Boot Order

The boot order command **bo** is used to set a default ASN-9000 booting order. The syntax of the **bo** command is:

```
bo set <value> bo [show] bo unset
```

<b>show set unset</b>	Sets, shows or unsets the boot order. The boot order designates the order of sources from which the ASN-9000 attempts to boot.
<b>m n</b>	The boot values can be the Flash Memory Module ( <b>m</b> ) and/or network ( <b>n</b> ). The boot order can be set in any order.

If more than one boot source is specified, the ASN-9000 attempts them in the order specified. For example: if **mn** is entered, the ASN-9000 attempts to boot from Flash Memory and then the network.



If more than one boot source is specified, the configuration files on each source should match to prevent an erroneous configuration from being loaded into the ASN-9000.

The following examples show the results of the various boot order command options:

```
264:ASN-9000:nvram# bo
bo          m  (flash-module)
265:ASN-9000:nvram# bo set mn
266:ASN-9000:nvram# bo
bo          mn  (flash-module,net)
267:ASN-9000:nvram# bo unset
268:ASN-9000:nvram# bo
bo          (not set, defaults to "f")
270:ASN-9000:nvram# bo set m
```

## 7.1.2 My IP Address

The **myip** command is used to set the IP address of the ASN-9000. The syntax of this command is:

```
myip [show] myip set <ipaddr> myip unset
```

<b>show set unset</b>	Specifies whether to show, set, or unset (clear) the IP address associated with the ASN-9000.
-----------------------	---

**<ipaddress>** Specifies the IP address of the ASN-9000.

```
280:ASN-9000:nvram# myip
myip          169.144.86.54
281:ASN-9000:nvram# myip set 169.144.86.55
282:ASN-9000:nvram# myip unset
```

### 7.1.3 My Subnet Mask

Use the **mysm** command to set the subnet mask of the ASN-9000:

```
mysm [show] mysm set <ipaddr-mask> mysm unset
```

**show|set|unset** Specifies whether to show, set, or clear the IP subnet mask for the interface.

**<ipaddr-mask>** Specifies the IP subnet mask.

```
285:ASN-9000:nvram# mysm
mysm          255.255.255.0
286:ASN-9000:nvram# mysm set 255.255.255.255
287:ASN-9000:nvram# mysm unset
```

### 7.1.4 File Server IP Address

Use the **fsip** command to set the IP address of the file server:

```
fsip [show] fsip set <ipaddr> fsip unset
```

**show|set|unset** Specifies whether to show, set, or clear the IP address of the file server.

**<ipaddress>** A file server's IP address.

```
290:ASN-9000:nvram# fsip
fsip          (not set)
291:ASN-9000:nvram# fsip set 169.144.86.49
292:ASN-9000:nvram# fsip unset
```

### 7.1.5 Gateway IP Address

Use the **gwip** command to set the IP address of the gateway server:

```
gwip [show] gwip set <ipaddr> gwip unset
```

**show|set|unset** Specifies whether to show, set, or clear the IP address of the gateway router.

**<ipaddress>** Specifies an intervening router's (gateway's) IP address.

## 7.1.6 Crash Reboot

The **crashreboot** command instructs the ASN-9000 to reboot automatically following a system crash. The syntax of this command is:

```
crashreboot [show] crashreboot set crashreboot unset
```

**show/set/unset** Specifies whether the switch automatically attempts a reboot following an unexpected system crash. The default is **set**, which causes the ASN-9000 to attempt a reboot following a crash. We recommend that this setting not be changed unless instructed to do so by FORE Systems TAC.

```
302:ASN-9000:nvram# crashreboot
crashreboot      (set)
303:ASN-9000:nvram# crashreboot set
304:ASN-9000:nvram# crashreboot unset
```

## 7.1.7 Slot Segments

The **slotsegs** command is used to allocate segments to specific slots. The syntax of this command is:

```
slotsegs [show] slotsegs[<n>] [show] slotsegs[<n>] set
<segment-count> slotsegs[<n>] unset
```

**<slot>** Specifies the slot number for which segments are being allocated.

**<num>** Specifies the number of segments being allocated to the specified slot.



The brackets around the slot number must be entered as they are part of the command.

## 7.2 RIPv2 Authentication

RIPv2 supports encrypted packet transmission using the MD5 algorithm to authenticate route and table updates. The MD5 algorithm allows packets to be encrypted at a source ASN-9000 and decoded at a destination ASN-9000 containing the same encryption key and key-string (password). Because the keyID is not transmitted over the network, but is set at each end, it reduces the likelihood of a successful attack on the network.

MD5 authentication is only supported in RIPv2. It does not work in RIPv1. RIPv2 must be enabled and running on all interfaces that require authentication. Additionally, RIPv2 authentication is not supported on interfaces that are configured for both RIPv1 and RIPv2; interfaces must be configured for RIPv2 only.

The MD5 key must be set up on the ASN-9000s at both sides of the connected interfaces in order for the authentication to take place. The keyid and the key-string must be the same on both ASN-9000s. Refer to RFC-2082 for a discussion on RIPv2 authentication using the MD5 encryption algorithm. The syntax for the **md5key** command is:

```
md5key [show] md5key[<keyid>] [show] md5key[<keyid>]
      set <key-string> md5key[<keyid>] unset
```

<b>[keyid]</b>	Specifies the number or identifier of the MD5key. The number must be a whole number between 1 and 255. The brackets around the keyid are part of the command and must be included. This variable is not required with the <b>show</b> argument.
<b>&lt;key-string&gt;</b>	Specifies the password for the encryption. The maximum password length is 16 characters. This variable is only required with the <b>set</b> argument.

Example of the **md5key** command are contained below:

```
338:ASN-9000:nvram# md5key
Total keys reserved: 0
339:ASN-9000:nvram# md5key[1] set abcdefghijklmnopq
Error: Authentication string is too long: abcdefghijklmnopq
340:ASN-9000:nvram# md5key[1] set abcdefghijklmnop
341:ASN-9000:nvram# md5key
md5key[ 1]      (set)
                  Total keys reserved: 1
342:ASN-9000:nvram# md5key[1] unset
343:ASN-9000:nvram# md5key
                  Total keys reserved: 0
```

In this example, key 1 is set with the keyID (password) of “abcdefghijklmnop.” This is the only time the keyID displayed.



## CHAPTER 8

# Media Subsystem Commands

This chapter describes the *ForeRunner* ASN-9000 media subsystem commands. The media subsystem commands relate to the physical media and bridging configuration information. This chapter explains the commands that allow the following:

- Show bridging-related configuration information
- Clear, enable, disable or show Inter-Segment statistics
- Enable, disable or show port-by-port statistics
- Enable or disable transmission and reception of packets on given segments
- Set or show segment names for given segments
- Enable, disable or show automatic segment state detection for all or specified segments
- Set segment state detection thresholds for specified segments
- Show port-level status for UTP ports
- Show or clear media-level statistics for UTP ports or segments

## 8.1 Displaying Bridge-Related Configuration

---

The current port and segment configuration can be displayed using the **config** command. The **config** command is issued from the **media** subsystem and displays the following bridge-related information:

- Forwarding status of segments
- UTP port receiver enable/disable status
- Automatic segment state detection
- Segment names
- Port level statistics collection
- Inter-segment statistics collection

The syntax of the **config** command is:

```
config [show] [<params>] [<disp-restrictors>]
```

<b>&lt;params&gt;</b>	Specifies a comma separated list of
monitor	show lan-monitor status
segment	forwarding enabled or not
[port]receive	receivers enabled or not
ssd	segment-state detection
[segment]names	names assigned to segments
portstats	collection enabled or not
isstats	collection method enabled or not

<b>&lt;disp-restrictors&gt;</b>	Specifies the segment or segments to display status on in a segment list <seglist>.
---------------------------------	---

Entering **config 2.1**, to display the bridge-related configuration information of port 2.1, displays the following:

```
24:PHswitch:media# config 2.1
Port Monitoring
-----
Packets...
not being monitored on segment 2.1
Forwarding status of segments
-----
2.1 :enabled
UTP port receiver enable/disable status
-----
Slot 3: .
Slot 2: . . . . .
Slot 1: . . . . .
. . . . .
Automatic segment state detection
-----
Segment 2.1 : enabled (currently good)
Segment names
-----
2.1 : Port_17
Port level statistics collection: currently disabled.
Inter-Segment Statistics collection is disabled
25:PHswitch:media#
```

## 8.2 Inter-Segment Statistics

Inter-Segment Statistics can be cleared, enabled, disabled or displayed using the **isstats** command. The **isstats show** command displays the packet statistics of packets between segments of the installed ASN-9000 NIMs, if statistic collection is enabled. The syntax of **isstats** is:

```
isstats [show] [<params>] [<disp-restrictors>] isstats
clear|enable|disable
```

- <params>** Specifies a comma separated list of p,o for packets and/or octets.
- <disp-restrictors>** Specifies a fr[om]=<seglist> to=<seglist> list of segments.

The following display shows two variations of displaying inter-segment statistics for either a range of segments or specific segments. Note that a range is separated with a hyphen.

```
12:PHswitch:media# isstats show 1.1-1.4
Segment to segment statistics collection is disabled
FROM      TO>      1.1      1.2      1.3      1.4
1.1  :(pkts)      (0)      (0)      (0)      (0)
      :octets      0        0        0        0
1.2  :(pkts)      (0)      (0)      (0)      (0)
      :octets      0        0        0        0
1.3  :(pkts)      (0)      (0)      (0)      (0)
      :octets      0        0        0        0
1.4  :(pkts)      (0)      (0)      (0)      (0)
      :octets      0        0        0        0
14:PHswitch:media# isstats show 1.1,2.1
Segment to segment statistics collection is disabled
FROM      TO>      1.1      2.1
1.1  :(pkts)      (0)      (0)
      :octets      0        0
2.1  :(pkts)      (0)      (0)
      :octets      0        0
15:PHswitch:media#
```

## 8.3 Ethernet LED Modes

---

The **ledmode** command is not implemented on the ASN-9000.

## 8.4 Operating-Mode

---

The **operating mode** command is not implemented on the ASN-9000.

## 8.5 UTP Port Receiver Status

---

The **portreceive** command is not implemented on the ASN-9000.

## 8.6 Displaying Port-Level Statistics

---

The **monitor** command is not implemented on the ASN-9000.

## 8.7 Configuring Packet Forwarding on Segments

Under certain circumstances it is desirable to disable the transmission of packets and the reception of packets on a specific segment. To enable or disable packet forwarding and reception, use the **segment** command. Here is the syntax for the **segment** command:

```
segment penable|pdisable <segment-list>
```

**<segment-list>** Specifies a single segment, dash-separated segment range, or a comma-separated list of segments.

Here are some examples of this command:

```
195:ASN-9000:media# segment pdisable 2.21-2.28
Segment 2.21: disabled
Segment 2.22: disabled
Segment 2.23: disabled
Segment 2.24: disabled
Segment 2.25: disabled
Segment 2.26: disabled
Segment 2.27: disabled
Segment 2.28: disabled
196:ASN-9000:media#
```

In the above example, a range of segments, 2.21 to 2.28, is disabled. In the example below, a comma-separated list of segments is enabled. Note that there are no spaces in the comma-separated list.

```
198:ASN-9000:media# segment penable 1.1,1.3,1.5,1.6
Segment 1.1: enabled
Segment 1.3: enabled
Segment 1.5: enabled
Segment 1.6: enabled
199:ASN-9000:media#
```

## 8.8 Configuring Segment Names

When the ASN-9000 first boots and assigns segments, the segment numbers are assigned from bottom to top and are named starting with “Port\_1.” To rename or display the port names on the ASN-9000, issue the **segmentname** command. Here is the syntax for the **segmentname** command:

```
segmentname|name sset <name> <seglist> segmentname|name  
[show] [<seglist>]
```

<b>sset</b>	Sets the segment name for the specified segment(s).
<b>show</b>	Displays the segment name(s) for the specified segment(s).
<b>&lt;name&gt;</b>	Specifies the name to use as a replacement for the default “Port_x” name. This variable is not required when using the <b>show</b> argument.
<b>&lt;seglist&gt;</b>	Specifies the segment number of the segment to be renamed. This variable must be a single segment number when renaming a segment. This variable may be a dash-separated range or a comma-separated list when used with the <b>show</b> argument. When the <b>show</b> argument is used and the <b>&lt;seglist&gt;</b> variable is not used, all segments are displayed.

An example of the **segmentname** command is shown below:

```
49:PHswitch:media# segmentname sset tpubs 2.1
Segment 2.1 named: tpubs
51:PHswitch:media# segmentname show 2.1
Segment names:
2.1 : tpubs
52:PHswitch:media#
```

## 8.9 Automatic Segment-State Detection

---

Automatic segment-state detection recognizes is a segment is down and automatically disables bridging and routing on that segment. When it has been detected that a segment’s state has changed, the segment is disabled (taken out of service) and the software is marked to denote the change. The updated segment state is displayed when the **ssd** command is issued.



If automatic segment-state detection on a segment is disabled, the segment’s state is always reported as “good” and interface states are always reported as “up” in the software. For information about a segment’s or interface’s state, enable automatic segment-state detection for that segment.

An ATM segment is considered to be down if the ELAN or the physical link to the AMA goes down.

## 8.9.1 Software Behavior When Disabled

When a segment is disabled, no packets are bridged or routed on that segment. Bridging and routing do not occur whether the segment is disabled by automatic segment-state detection or by issuing the `segment pdisable` command. See Section 8.7 for information on the `segment pdisable` command.

## 8.9.2 Default Setting

The default setting for the automatic segment-state detection for an ATM segment is enabled.

## 8.10 Setting Segment-State Threshold

---

The `ssdthreshold|ssdt` command is not implemented in the ASN-9000.

## 8.11 Status

---

The `status` command is used to display the port-level status of the specified, or all, UTP ports on NIMs installed in the ASN-9000. As shown in the below example, the Link Test, Partitioning and Polarity of all UTP ports is displayed.

```
11:PHswitch:media# status
Link Test:
Slot 3: -
Slot 2: Y - - - - - - - - - - - - - - - - - -
Slot 1: - - - - - - - - - - - - - - - - - -
- - - - -
Partitioning:
Slot 3: .
Slot 2: . . . . .
Slot 1: . . . . .
. . . . .
Polarity:
Slot 3: .
Slot 2: . . . . .
Slot 1: . . . . .
. . . . .
12:PHswitch:media#
```

The syntax for the `status` command is:

```
status [show] [<params>] [<display-restrictors>]
```

<b>w&lt;params&gt;</b>	Specifies a comma-separated list of link, partition, polarity or all.
<b>&lt;disp-restrictors&gt;</b>	Specifies the segment or segments to display the port-level UTP port status.

## 8.12 Media-Level Statistics

---

The **stats** command is used to display media-level statistics. If port statistics are being collected (refer to Section 8.6) they are displayed in place of segment-wide statistics unless the **-s** flag is used. Segment-level statistics can be cleared but not disabled. When segment-level statistics are cleared, they are reset to zero and immediately begin to increment as packets are received and sent by the switch. The syntax of the **stats** command is:

```
stats [show] [-p|-s] [<params>] [<display-restrictor>]  
stats clear
```

<b>-p</b>	Display port statistics, if available.
<b>-s</b>	Display segment-level statistics.
<b>&lt;params&gt;</b>	Specifies a comma-separated list of the following (or "all")  pi, po, oi, oo, bpi, bpo, pu, rbe, xbe, fcs, fa, c, rc, tc, q, gp, cu, lc, er, tm  In addition to above parameters, the following apply specifically to 2x8 Fast Ethernet Repeaters ce, drm, jbrs, se, rts, lcs, ap, iss, sfc, lsa  See Table 8.1 for parameter definitions.

**Table 8.1 - Segment Level Statistic Parameters**

Parameter	Description
pi	packets in
po	packets out
oi	octets in
oo	octets out
bpi	broadcast packets in
bpo	broadcast packets out
pu	peak utilization
rbe	receive buffer errors
xbe	transmit buffer errors
fcs	frame check sequence errors
fa	frame alignment errors
c	segment collisions
rc	port collisions
tc	transmit collisions
q	output queue length
gp	giant packets
cu	current utilization
lc	local carrier
er	excessive retries
tm	table miss
ce	coding errors
drm	data rate mismatch
jbrs	jabbers
se	short events
rts	runts
lcs	rate collisions
ap	auto partitions
iss	isolates
sfc	source address field changes
lsa	source address of the last incoming packet.

## *Media Subsystem Commands*

- |                                   |  |
|-----------------------------------|--|
| <b>&lt;display restrictor&gt;</b> | Specify a specific segment, range of segments or a comma-separated list of segments. |
| <b>clear</b>                      | Clears all port and segment-level statistics.  |

The **stats show** command displays media-level statistics. If portstats are being collected, portstats statistics are displayed in place of segment-wide statistics unless the -s flag is used.

# CHAPTER 9

## Host Subsystem Commands

This chapter describes the commands in the `host` subsystem and tells you how to use these commands to perform the following tasks:

- Display the TCP configuration settings.
- Display the TCP table.
- Display TCP, TELNET, and UDP statistics.
- Clear TCP, TELNET, and UDP statistics.
- Set the connection time.
- Set the keep-alive interval.
- Kill a TCP connection.
- Display the UDP table.

The ASN-9000 software's `host` subsystem includes an implementation of the TCP (Transmission Control Protocol) stack, a connection-oriented, industry-standard protocol for moving data between nodes in a network environment. In particular, TCP is used by TELNET, a program that allows workstations to communicate with the ASN-9000 using an in-band network connection.

To define TCP filters for controlling access to your network, refer to the *ForeRunner ASN-9000 Filters Reference Manual*.

### 9.1 Accessing the Host Subsystem

---

To access the `host` subsystem, issue the following command at the runtime command prompt:

```
host
```

To list available commands in the `host` subsystem issue `help` or `?`.

## 9.2 Displaying the TCP Configuration

---

Use the **config [show] tcp** command to display configuration parameters used by the host subsystem.

```
67:ASN-9000:host# config show tcp
TCP Configuration
-----
```

```
Round Trip Algorithm:          vanj
Min Rexmit Interval:          1000 ms
Max Rexmit Interval:          64000 ms
Max Connections Allowed:      2

connection-idle-time:          20 minutes
keep-alive-interval [kainterval]: 75 seconds
keep-alive delay [kadelay]:    1200 seconds
Time to disconnect on idle conn: 30 minutes 0 seconds
```

This display shows the following information about the current TCP configuration parameters:

- The round-trip algorithm used by the ASN-9000 software is the Van Jacobson algorithm.
- The minimum retransmit interval is 1,000 milliseconds.
- The maximum retransmit interval is 64,000 milliseconds.
- The maximum number of simultaneous TELNET (TCP) connections that can be supported is two.
- The connection-idle time is 20 minutes.
- The keep-alive interval is 75 seconds.
- The keep-alive delay is 1200 seconds.
- The time allowed before an idle connection is automatically disconnected is 30 minutes. This value is based on the values of the connection-idle time and the keep-alive interval.

### 9.3 Displaying the TCP Table

The *TCP table* lists the active TCP connections between the ASN-9000 and other devices. Use the **status show tcp** command to list the TCP table. Here is an example of the TCP table:

```
7:ASN-9000:host# status show tcp
Active TCP Connections
Conn Id  Rem IP Addr  Rem Port  Loc IP Addr  Loc Port  Conn. State
-----
16       147.128.128.128  1043      147.128.128.64  23        ESTABLISHED  **
17       147.128.128.8   1201      147.128.128.64  23        ESTABLISHED
List of registered UDP clients:
161 SNMP
520 RIP
```

For each TCP connection to the ASN-9000, the TCP table shows information under the following headings:

Conn ID	A unique integer that identifies the connection. This identifier can be used to terminate the connection using the <b>kill &lt;connection-id&gt;</b> command.
Rem IP Addr	The IP address of the remote device that initiated the connection.
Rem Port	A process port number for the remote device (management station). Note that the process port number is unrelated to the ASN-9000 physical port or segment numbers. It is assigned by the remote operating system.
Loc IP Addr	The IP address of the local device. This is always the ASN-9000.
Loc Port	A process port number for the ASN-9000. This is unrelated to the ASN-9000 physical port or segment numbers. It is a “well-known” port number used by the TELNET process.
Conn. State	The connection state is one of the following states of the standard TCP software state machine: <div><div>CLOSED CLOSE-WAIT FIN-WAIT-1 LAST-ACK SYN-RECEIVED</div><div>CLOSING ESTABLISHED FIN-WAIT-2 LISTEN SYN-SENT TIME-WAIT</div></div>

Most of these states are never displayed by the **status show tcp** command because they occur for a short time. Connections in the CLOSED or LISTEN state are not displayed.

The current TELNET session (if you are connected through TELNET) is indicated by two asterisks (\*\*) following the table entry for that session.

## 9.4 Displaying Statistics

---

The **host** subsystem maintains statistics on TCP, TELNET, and UDP packets. TCP and UDP statistics are a superset of the corresponding statistics provided in the SNMP MIB. (There is no TELNET MIB.)

The software maintains two copies of each TCP, TELNET, and UDP statistics counter:

- Count since last statistics clear.
- Count since last system reset.

Use the **stats** command to display statistics. Here is the syntax for this command:

```
stats [show|clear] [-i] [-t] tcp|tel[net]|udp
```

- |                       |   |
|-----------------------|---|
| <b>tcp telnet udp</b> | Specifies the type of protocol for which you want to display packet statistics.   |
| <b>-t</b>             | Displays statistics totals collected since the last system reset, rather than the statistics collected since the last statistics clear. |

In this example, TCP statistics collected since the last statistics clear are displayed.

```
62:ASN-9000:host# stats tcp
```

```
TCP Connection & Pkt statistics (count since last stats clear):
Active Opens:                0
Passive Opens:               5
Failed Conn Attempts:        0
Resets In Estb State:         0
Current Open Conns:          2
Segments Received:           1088
Segments Sent:               1030
Rexmitted segments:          2
Segments Rcvd With Err:      0
Resets Sent:                  0
Short Segments Rcvd:         0
```

## 9.4.1 Clearing Statistics

Use the **stats clear** command to clear statistics for TCP, TELNET, or UDP packets. Here is the syntax for this command:

```
stats clear [-i] [-t] tcp|tel[net]|udp
```

**tcp|tel[net]|udp** Specifies the type of protocol for which you want to clear packet statistics.

When you clear statistics, the counters that record statistics since the last clear are reset to zero. The *ForeRunner* ASN-9000 software immediately begins collecting new statistics. The counters that record statistics since the last system reset are unaffected by the **stats clear** command.

## 9.5 Setting TCP Session Parameters

---

The *ForeRunner* ASN-9000 software can kill idle TCP (TELNET) connections automatically using the following TCP configuration parameters:

- Connection-idle time (default 20 minutes).
- Keep-alive interval (default 75 seconds).

The combination of the two parameters above determines the time allowed to pass before an idle connection is automatically disconnected.

### 9.5.1 The Connection Idle Time and Keep Alive Interval

When a new TELNET connection is established, an idle timer is started. The idle timer is reset to 0 and restarted whenever there is activity on the connection. If the idle timer reaches a pre-set value, the *connection idle time*, then the ASN-9000 sends a keep-alive packet to the remote device. If there is still no activity, the ASN-9000 continues to send keep-alive packets at an interval called the *keep-alive interval*, until eight keep-alive packets are sent. If there is still no activity, then the connection is dropped.

A connection is automatically dropped if it is idle for a period of time equal to the connection-idle time plus eight times the keep-alive interval. Using the default values of these parameters, the maximum idle time is 30 minutes.

In addition, you can display and set the control characters used for displaying and editing text during a TELNET session.

### 9.5.1.1 Setting the Connection Idle Time

Use the **set kadelay** command to specify how long a TELNET connection can remain idle before the ASN-9000 sends keep-alive packets. Here is the syntax for this command:

```
set kadelay|kad <time>
```

**<time>** Specifies how many minutes the ASN-9000 allows a TCP (TELNET) connection to remain idle before sending keep-alive packets. The range for this value is 5 to 30 minutes; the default is 20 minutes.

### 9.5.1.2 Setting the Keep-alive Interval

Use the **set kainterval** command to specify how often the ASN-9000 sends keep-alive packets before ending a connection. Here is the syntax for this command:

```
set kainterval|kai <time>
```

**<time>** Specifies how often the ASN-9000 sends keep-alive packets before ending a connection. The range for this value is 30 to 240 seconds; the default is 75 seconds.

## 9.5.2 Closing a TCP Connection

At any time, you can end a TCP (TELNET) connection. The command you use depends upon whether you are ending the:

- Connection from which you are working.
- Connection other than the one from which you are working.

### 9.5.2.1 Current TCP Connection

Use the **logout** command to end the current TCP connection.

### 9.5.2.2 Another TCP Connection

Use the **kill** command to end a TCP connection other than the one in which you are working. You must issue this command from a session other than the one you are ending. Here is the syntax for this command:

```
kill <connection-id>
```

**<connection-id>** Specifies the ID assigned to the session by the ASN-9000 when the session was established. To determine what the connection ID is, use the **stats tcp** command to display the TCP table. The connection IDs are listed in the first column, under Conn ID.

## 9.6 Displaying the UDP Table

The ASN-9000 software contains agents that can respond to certain “well-known” UDP port requests, such as RIP packets and SNMP requests. The ports listed in the UDP port table are the port numbers that UDP clients register with the UDP protocol code.

When the ASN-9000 receives a UDP packet, it checks the UDP port number specified in the packet against the list of UDP ports in the UDP port table. If the ASN-9000 can respond to the UDP request, it does so. If the ASN-9000 cannot respond to the UDP request, it does one of the following:

- Drops the packet.
- Forwards the packet to the device specified by the IP Helper address, if an IP Helper address has been configured for the segment on which the UDP packet is received. See *Chapter 13, IP Subsystem Commands* for information on using the ASN-9000 IP Helper feature.

To display a complete list of the UDP protocol ports supported by the ASN-9000, issue the **status udp** command. The numbers and names are “well-known” UDP protocol port numbers and names as defined in RFC 1700. Here is an example of the display produced by this command:

```
81:ASN-9000:host# status udp
List of registered UDP clients:
 161 SNMP
 520 RIP
 67  BOOTPS
 68  BOOTPC
```

The UDP ports listed in this display indicate that the ASN-9000 contains agents for processing UDP packets sent to UDP protocol ports 161, 520, 67, and 68. In other words, the ASN-9000 supports the following types of UDP packets:

- SNMP
- IP RIP
- BOOTP (on server side)
- BOOTP (on client side)



For information on using bridge filters and templates, see the *ForeRunner ASN-9000 Filters Reference Manual*.

# CHAPTER 10 TFTP Subsystem Commands

The `tftp` subsystem contains the ASN-9000 implementation of TFTP (Trivial File-Transfer Protocol). Use the `tftp` subsystem commands to perform the following tasks:

- Set the default TFTP server.
- Display the default TFTP server.
- “Unset” the default TFTP server.
- Download or display a file stored on a TFTP server.
- Upload a file from the floppy drive or Flash Memory Module to a TFTP server.
- Load (activate) a configuration file stored on a TFTP server.
- Save ASN-9000 configuration changes to a configuration file on a TFTP server.

To use the commands in this subsystem, you must configure your TFTP server to support TFTP file transfers. The procedures for configuring your server depend upon the particular type of server you are using. See your server documentation for configuration information.

Also, the ASN-9000 segment that connects the ASN-9000 to the TFTP server must have an IP interface defined on it. For information about adding an IP interface, see *Chapter 13, IP Subsystem Commands*.

## NOTE

The TFTP protocol provides no authentication for any services, including downloading or changing files stored on the TFTP server. If you configure your TFTP server to allow the `tftp` commands to be used, anyone with access to the server can download or change files.

## 10.1 Accessing the TFTP Subsystem

To access the `tftp` subsystem, issue the following command at the runtime command prompt:

```
tftp
```

## 10.2 Considerations

---

ASN-9000 TFTP commands work with many types of TFTP servers, including servers running UNIX, DOS, Windows, OS/2, or other operating systems. The following considerations apply to TFTP servers that are running UNIX, a very common platform for TFTP.

Regardless of the platform used for the TFTP servers in your network, we recommend that you consult your server's documentation regarding:

- File permissions (not applicable to some operating systems).
- Conventions for pathnames and file names.

If you experience problems uploading or downloading files between the ASN-9000 switch and your TFTP server, you often can resolve the problems by verifying whether the switch has or needs read and write access to the server, and how file names need to be specified to the switch.

### 10.2.1 TFTP Commands and UNIX Read/Write Permissions

To use the ASN-9000 TFTP commands to upload or download a file, or to save or read a ASN-9000 configuration file, the proper UNIX read/write permissions must be set on the TFTP server. On most servers, permissions are controlled separately for users, groups, and "others." The TFTP server considers the ASN-9000 system to be among the "others."

You can control read/write access to ASN-9000 files and directories on the TFTP server by setting the read and write permissions. On most UNIX systems, you can display permissions information using the UNIX **ls** command. Here is an example of the permissions information displayed for a file on a TFTP server. The display on your TFTP server might be different.

```
$ ls -l
total 3
-rw-rw---- 1 mrspat      622 Jul 19 15:09  Lab1.env
-rw-rw-r-- 1 ethan       643 Jul 19 15:11  Lab2.env
-rw-rw--w- 1 sascha      611 Jul 19 15:13  Lab3.env
-rw-rw-rw- 1 stripie     698 Jul 19 15:15  Lab4.env
-rw-rw-rw- 1 tiger      698 Jul 19 15:15  Lab5.env
```

The text shown in bold is the permission information for each file, for "others."

- In this example, no read or write permissions are enabled for others for Lab1.env. Consequently, you cannot download this file or upload a file by this name using the ASN-9000 TFTP commands.
- Read permission, but not write permission, is granted to the file Lab2.env for others. You can use the ASN-9000 TFTP commands to display or download this file, but you cannot upload a file by this name.

- The file `Lab3.env` cannot be downloaded using the ASN-9000 TFTP commands. You can, however, upload a file by this name.
- Finally, both read and write permissions are enabled for the file `Lab4.env` and `Lab5.env`. You can download these files and upload a file by these names.

On most TFTP servers, you can change permissions using the UNIX `chmod` command. See the documentation for your UNIX shell for details.

## 10.2.2 PathNames

Depending upon your TFTP server configuration, you might need to specify pathnames with the TFTP commands.

On some TFTP servers, when you use the ASN-9000 TFTP commands to upload or download files, the ASN-9000 software understands file names according to where the ASN-9000 system accesses the server. You can upload or download only those files that are located in the directory where the ASN-9000 accesses the server, or in one of that directory's subdirectories. Also, if the file is in a subdirectory, you must specify the pathname along with the file name.

For example, suppose you configure your TFTP server to allow the ASN-9000 system to access the server at a directory called TFTP.

```
TFTP
    fore
        ph
            ethan.env
            sascha.env
```

All directories below the TFTP directory are considered part of the pathname for the files stored there. Relative to the ASN-9000, the pathname for the files `ethan.env` and `sascha.env` is `fore/ph`. To download the file `ethan.env`, you would issue the following command:

```
get -a fore/ph/ethan.env ethan.env
```

<b>-a</b>	Specifies net-ASCII mode. (Files are transferred in binary mode by default.)
<b>fore/ph/ethan.env</b>	Is the file name, including the pathname.
<b>ethan.env</b>	Is the name you want the file to have on the ASN-9000. This name must be in DOS format (filename.ext).

## 10.2.3 File Naming Conventions

Local file names are optional with the ASN-9000 `tftp get` command. You can omit the local file name if the file is not in a subdirectory and the file name is eight characters or fewer in length with an extension no longer than three characters.

Suppose the ASN-9000 has access to the TFTP server at the TFTP directory, as in the following example:

```
TFTP
    fore
        ph
            sascha.env
            ethan.env
            lotsofdots
```

If the `get` command is issued without specifying the local file name (`sascha.env`), an error message is displayed on the ASN-9000 terminal.

In addition, the ASN-9000 uses DOS file-naming conventions, but the file `lotsofdots` does not fit the DOS file-naming conventions. To download `lotsofdots`, you need to specify a local file name that fits the DOS file naming conventions, as in the following example:

```
get -a fore/ph/lotsofdots spots
```

In this example, the file `lotsofdots` is named `spots` on the ASN-9000.

## 10.2.4 Remote File Names

Some TFTP servers require that the remote file name exist on the server before you can write to that file name. If your server requires that the file name already exist, create a zero-length file on the server, then specify the name of that file as the remote file name with the `put` or `savecfg` commands.

Also, on some TFTP servers, files that you overwrite on the server are not properly truncated. When you overwrite an existing file on the TFTP server, if the older version of the file is longer than the new file, the older version is not truncated properly by the server. As a result, the new version of the file contains part of the older version of the file. If you are unsure whether the new version will completely replace the older version of a file, do one of the following:

- Remove the older version of the file, then save the new version.
- If your server requires that the file name be present on the server before you can copy to it, create a zero-length file under a new name, then save the ASN-9000 file under the new name. After the new file is copied to the server, delete the older version of the file and rename the new file as desired.

## 10.3 Setting, Displaying, or Unsetting the Default Server

The commands described in the following sections let you specify a particular TFTP server to be used in the file operations described in this chapter.

If you choose not to specify a default TFTP server, you still can specify a server with the individual TFTP commands.

### 10.3.1 Setting the Default TFTP Server

Use the **set** command to specify the default TFTP server. Here is the syntax for this command:

```
set server <ipaddress>
```

**<ipaddress>** Specifies the IP address of the TFTP server you want to use as the default. Specify the address in dotted-decimal notation.

You can only have one active TFTP server at a time. Setting a new default TFTP server's IP address replaces the existing TFTP server's IP address.

### 10.3.2 Displaying the Default TFTP Server

Use the **show** command to display the IP address of the default TFTP server. Here is the syntax for this command:

```
show server
```

The **show server** command shows you the TFTP server at the IP address you specify with the **set server** command.

### 10.3.3 Removing the Default TFTP Server Setting

Use the **unset** command to remove the default TFTP server setting. Here is the syntax for this command:

```
unset server
```

## 10.4 Downloading or Displaying a File

---

Use the **get** command to display or download a file stored on a server. Here is the syntax for this command:

```
get [-h <host>] [-a] <remote-file> [<local-file>|tty]
```

**-h <host>** Specifies the IP address, in dotted-decimal notation, of the TFTP server. If you do not specify this argument, the default server is used. The default server is specified using the **set server** command. (See Section 10.3.1.)

**-a** Forces the transfer to take place in net-ASCII transfer mode, rather than octet mode. Octet mode transfers the file, including end-of-line characters, exactly as it is stored on the server. Net-ASCII changes the end-of-line characters to be compatible with the display or storage device that receives the file.

Use the default (octet-mode) to download software image files (ex: 7f, 7pe, 7atm). Use the net-ASCII mode to download configuration files, environment files, and other text files.

If you plan to display the file on your management terminal (by specifying **tty** as the local file name), omit this argument. The file is automatically transferred in net-ASCII format.

**<remote-file>** Specifies the name of the remote file. Specify the name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called **transfer** and this directory is specified as the TFTP home directory, do not specify **transfer** as part of the file name.

**<local-file>|tty** If you omit this argument, the ASN-9000 assumes that you want to use the same file name on the server and include the pathname (if any) in the file name.

If you omit the **<local-file>** argument or specify a local file name, the file is written to a local storage device.

If the file name on the server is an invalid pathname on the ASN-9000, an error message is displayed on the ASN-9000 terminal.

If you specify **tty**, the file is not downloaded to your system, but an image of the file is displayed on the management terminal. You can display the file from within a TTY (RS-232) session or a TELNET session.

If you do not specify a TFTP server name and no default server name has been configured, an error message is displayed. To configure a default server name, use the **set server** command. (See Section 10.3.1.)

## 10.5 Uploading a File

Use the **put** command to upload a file stored on the floppy diskette or Flash Memory Module of a ASN-9000 to a TFTP server. Here is the syntax for this command:

```
put [-h <host>] [-a] <localfile> [<remote-file>]
```

**-h <host>** Specifies the IP address, in dotted-decimal notation, of the TFTP server. If you do not specify this argument, the default TFTP server is used. The default TFTP server is specified using the **set server** command. (See Section 10.3.1.)

**-a** Forces the transfer to take place in net-ASCII transfer mode, rather than octet mode. Octet mode transfers the file, including end-of-line characters, exactly as it is stored on the server. Net-ASCII changes the end-of-line characters to be compatible with the display or storage device that receives the file.

Use the default (octet-mode) to download software image files (ex: 7PE, 7ATM). Use the net-ASCII mode to download configuration files, environment files, and other text files.

**<local-file>** Specifies the local file name.

**<remote-file>** Specifies the name of the file as you want it to appear on the server. Specify the name that is meaningful to the TFTP program on the server. For example, if the name with the path of the server contains a

subdirectory called `transfer` and this directory is specified as the TFTP home directory, do not specify `transfer` as part of the file name.

Here is an example of the use of this command. In this example, an environment file is uploaded to a TFTP program on a UNIX server.

```
12:ASN-9000:tftp# put -a /fore/env/ASN-90001.env
177.177.45.20:/fore/env: 833 bytes
13:ASN-9000:tftp#
```

Notice that a pathname is specified with the file name in this example. Make sure you specify the pathname that is meaningful to your TFTP program.

On UNIX machines, if the write permission for “others” is not enabled on the TFTP server for the file name or the directory to which you are trying to write the file, a message such as the following is displayed:

```
12:ASN-9000:tftp# put -a ASN-90001.env
tftpWrite: Peer generated error
tftp: Permission denied: Access violation
13:ASN-9000:tftp#
```

If you receive this error, check the file and directory permissions for “others” on the TFTP server.

If your TFTP program is on a UNIX machine and that machine requires that the file name already exist, but the file does not yet exist on the server, a message such as the following is displayed on the ASN-9000 terminal:

```
14:ASN-9000:tftp# put ASN-90002.env
tftpWrite: Peer generated error
tftp: File not found: File not found
15:ASN-9000:tftp#
```

## 10.6 Loading a Configuration File

---

During normal run-time operation of the ASN-9000, you can read (load) a configuration file stored on a remote TFTP server. To do so, issue the following command:

```
readcfg|rdcfg [-v] [-h <host>] <remote-file>
```

**-v** Displays commands to the open TELNET session as they are executed.

- h <host>** Specifies the IP address of the TFTP server. If you do not specify this argument, the default TFTP server is used. (The default TFTP server is specified using the **set server** command. See Section 10.3.1.
- <remote-file>** Specifies the name of the configuration file you want to load. Specify the name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called `configs` and this directory is specified as the TFTP home directory, do not specify `configs` as part of the file name.

As with the **get** command, if you do not specify a host server name and no default server name has been configured, an error message is displayed.

## 10.7 Saving a Configuration File

During normal run-time operation of the ASN-9000, you can save the ASN-9000's current configuration to a file on a remote TFTP server. To save configuration files, issue the following command:

```
savecfg [-h <host>] <remote-file>
```

- h <host>** Specifies the IP address of the TFTP server. (The default server is specified using the **set server** command. See Section 10.3.1
- <remote-file>** Specifies the name of the configuration file you want to save. Specify the name that is meaningful to the TFTP program on the server. For example, if the server contains a subdirectory called `configs` and this directory is specified as the TFTP home directory, do not specify `configs` as part of the file name.

On UNIX-based TFTP servers, if the write permission for “others” is not enabled for the configuration file name or the directory to which you are trying to write the file, a message such as the following is displayed:

```
16:ASN-9000:tftp# savecfg ace.cfg
tftpWrite: Peer generated error
tftp: Permission denied: Access violation
17:ASN-9000:tftp#
```

If you receive this error, check the file and directory permissions for “others” on the TFTP server.

## *TFTP Subsystem Commands*

If your UNIX-based server requires that the file name already exist, but the file does not yet exist on the server, a message such as the following is displayed on the *ForeRunner* ASN-9000 terminal:

```
18:ASN-9000:tftp# savecfg ace.cfg
tftpWrite: Peer generated error
tftp: File not found: File not found
19:ASN-9000:tftp#
```

# CHAPTER 11

## Bridge Subsystem Commands

The ASN-9000 software contains implementations of IEEE 802.1d bridging and the 802.1d Spanning-Tree protocol. This chapter describes the bridge subsystem commands you can use to perform the following tasks:

- Show the bridge configuration
- Show and manage the bridge table (includes changing the aging interval for dynamic (learned) entries)
- Show, add, and delete bridge groups
- Show the bridging status of an ASN-9000 segment
- Enabling, disabling, and configuring Spanning-Tree
- Display or clear packet, bridge, and segment statistics
- Display and clear the bridge cache

### 11.1 Accessing the Bridge Subsystem

---

To access the `bridge` subsystem, issue the following command from any command prompt:

```
bridge
```

### 11.2 Showing the Bridging Configuration

---

Use the `config` command to display the bridge configuration parameters. The syntax for this command is:

```
config [show] [<argument-list>]
```

**<argument-list>** Specifies the configuration parameters you want to display. You can specify an argument, a comma-separated list of arguments, or `all` for all arguments. Table 11.1 lists the arguments you can specify. The default is `all`.

**Table 11.1 - config Command Arguments**

command arguments	command descriptions
vars	The aging time for entries in the bridge table. This also shows if learning is enabled.
groups	The currently defined network groups.
templates	All logical filtering templates that are defined.
rules	All logical filtering rules that are defined.
<b>filters</b>	The packet-forwarding restrictions for all segments. This includes the source and destination logical filtering rules and whether or not learned entries are blocked.
spantree   st	All parameters that are configured for the Spanning-Tree Algorithm.



Command syntax for the **lrule** and **template** commands, are located in the *ForeRunner ASN-9000 Filters Reference Manual*.

The following example shows the type of information displayed by the **show config** command when issued without arguments.

```
46:ASN-9000:bridge# show config
Spanning Tree
  Status:          Enabled
  System Priority:  8000
  Spanning Tree Add: 01-80-c2-00-00-00
  My Bridge Address: 00-00-ef-02-42-50
  Max Age:          Curr val: 21  (Config val: 21)
  Hello Time:       Curr val: 4   (Config val: 4)
  Forward Delay:    Curr val: 16  (Config val: 16)
  Send Fast Hellos: Disabled
  Fast Hello Params: HelloTime:1sec, HighUtil:70%, LowUtil:50%
  Seg Prio PathCost DesignatedBridge DesSeg DesCost StaChngs
  ---
1.1 128 100      this-bridge 1      10      1
1.2 128 100      this-bridge 2      10      1
1.3 128 100      this-bridge 3      10      1
1.4 128 100      this-bridge 4      10      1
1.5 128 100      this-bridge 5      10      1
1.6 128 100      this-bridge 6      10      1
2.1 128 100      this-bridge 7      10      1
2.2 128 100      this-bridge 8      10      1
2.3 128 100      this-bridge 9      10      1
2.4 128 100      this-bridge 9      10      1
Bridge learning
  segment 1.1: on
  segment 1.2: on
  segment 1.3: on
  segment 1.4: on
  segment 1.5: on
  segment 1.6: on
  segment 2.1: on
  segment 2.2: on
  segment 2.3: on
  segment 2.4: on
Bridge table aging time: 60 minutes
IP bridging: disabled
Bridge Groups
Name          Segment List
-----
default       1.1, 1.2, 1.3, 1.4, 1.5
              1.6, 2.1, 2.2, 2.3, 2.4
```

## 11.3 Using the Bridge Table

---

The *bridge table* contains information about devices attached to the ASN-9000. The ASN-9000 uses the entries in the bridge table to bridge packets. Entries are added to the table automatically or manually.

- Entries Added Automatically

Each time the ASN-9000 bridging engine receives a packet, it checks the packet's source address against the MAC addresses listed in the bridge table. If the address is not listed in the table, the switch adds an entry to the table. The entry contains the source device's MAC address, the segment number on which the switch received the packet, and other information used for bridging.

- Entries Added Manually

You can create a static entry using the **bt add** command. A static entry is manually added to the bridge table, rather than learned by the bridge table. Static entries are not subject to aging and remain in the bridge table until you remove them. Moreover, they are saved in the configuration file when you save the file.<sup>1</sup> Use static entries when you want to ensure that the ASN-9000 always reaches a specific node on the same segment, or to configure a ASN-9000 connection to a multi-homed host.

### 11.3.1 Displaying the Bridge Table

To display the bridge table, issue this command:

```
bt [show] [-h] [-m] [-t] [<seglist>] [<ethaddr/ethpat>]
```

<b>&lt;seglist&gt;</b>	Specifies the segment(s) for which you want to display bridge table entries. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.  If you specify <b>all</b> , the entire bridge table is displayed. The default is <b>all</b> .
<b>&lt;ethaddr&gt; ethpat</b>	Is the MAC-layer hardware address of the device for which you want to display the ASN-9000 bridge-table entries. Specify the address as six hyphen-separated two-digit hexadecimal octets (ex: <b>08-00-20-0f-a5-ab</b> ). You can use <b>*</b> (asterisk) as a wildcard character in place of any of the octets.

---

<sup>1</sup>. To save an ASN-9000 configuration, issue the **system savecfg <file-name>** command.

- t** Displays the total number of entries in the table. The total is comprised of the total number of learned entries and permanent (static) entries. This argument also shows how many entries remain available in the bridge pool; that is, the number of entries for which the table still has room.
- h** Displays the hash displacements for the specified entries.
- m** Displays entries for multi-homed hosts.

**NOTE**

Because the **-h** and **-m** options display specific entries in the bridge table, they cannot be used with the **-t** option which displays total bridge entries.

Here is an example of the bridge table:

```
40:ASN-9000:bridge# bt show
```

```
Bridging table (aging time = 60 minutes)
```

Ethernet-address	Seg	Rule	Flags
00-00-ef-01-93-40	10	none	system permanent
01-80-c2-00-00-00	--	none	permanent bmast
aa-00-04-00-16-08	11	none	permanent bmcast
07-00-2f-e4-b3-ee	--	none	permanent bmcast
ff-ff-ff-ff-ff-ff	--	none	permanent bmcast

```
Total entries: 5, Learned entries: 0, Permanent Entries: 5
```

The bridge table contains the following information for each entry:

<b>Ethernet-address</b>	The MAC-layer hardware address of the device.
<b>Seg (Segment)</b>	The number of the segment to which the network joining the device to the ASN-9000 is attached. If the MAC-layer hardware address belongs to a multi-homed host, the segment number is shown as MH.
<b>Rule</b>	The number of a logical filtering rule applied to packets that are forwarded to or from this address. See the <i>ForeRunner ASN-9000 Filters Reference Manual</i> for information about defining rules.

**Flags** The software maintains certain flags in order to use and manage addresses in the bridge table. For example, switch entries such as the ASN-9000's own address are marked, and entries that haven't been used recently are flagged for possible deletion (aging).

Each entry in the bridge table can have one or more of the following flags:

bmcast	A broadcast/multicast address.
permanent	Most often, this flag indicates that the address is a static entry (created using the <code>bt add</code> command). Otherwise, it is a switch-defined entry.
spanning-tree	The industry-standard (IEEE 802.1d) multicast address used by the Spanning-Tree algorithm.
system	The factory-configured MAC-layer hardware address of the ASN-9000.
blank	In a typical application, most entries in the bridge table have none of the preceding flags set. Such entries are learned addresses that have been seen at least once since the last time the bridge table was aged.

### 11.3.2 Clearing the Bridge Table

Periodically, learned entries are automatically removed from the bridge table through aging. However, you can clear all learned entries from the table using this command:

```
bt clear
```

### 11.3.3 Adding an Entry to the Bridge Table

Use the `bt add` command to add a static (permanent) entry to the bridge table. The entry is added to the table as soon as you issue the command and remains in the table until you remove the entry. This command is helpful because adding static bridge entries is an effective

way to ensure that a ASN-9000 can always recognize a specific node that is permanently located on a segment. Unlike learned entries, static entries are not subject to aging. Here is the syntax for this command:

**bt add [<ethaddr>] [<seglist>]**

**<ethaddr>** Specifies the MAC address of the device.

**<seglist>** Specifies the segment(s) associated with the specified MAC address. To ensure that packets destined for the device are forwarded successfully, make sure you specify the segments to which the device is attached.



If you specify more than one segment, each segment is considered to be attached to a multi-homed host, and the flag M appears in the Segment column in place of a segment number.

Here is an example of how to create a bridge table entry for a multi-homed host. In this example, the entry is defined, then the bridge table containing the newly created entry is displayed. The entry is highlighted in bold type.

```
40:ASN-9000:bridge# bt add 00-00-EF-02-00-F0 1.1, 1.2, 1.3
address 08-00-20-00-ef-2b: added on 1.1, 1.2, 1.3
```

```
Bridging table (aging time = 60 minutes)
  Ethernet-address  Seg  Rule  Flags
  00-00-ef-01-93-40  10  none  permanent bmcast
  01-80-c2-00-00-00  --   none  system permanent
  aa-00-04-00-16-08  11  none  permanent bmcast
  07-00-2f-e4-b3-ee  --   none  permanent bmcast
  ab-00-00-03-00-00  --   none  permanent bmcast
  00-00-ef-01-10-20  10  none  system permanent
  aa-00-04-00-f1-33  10  none  system permanent
  ff-ff-ff-ff-ff-ff  --   none  permanent bmcast
```

```
Total entries: 9, Learned entries: 0, Permanent Entries: 9
```

### 11.3.4 Enabling and Disabling Bridge Learning

When bridge learning is enabled, MAC addresses from packets received on the ASN-9000 are recorded. The ASN-9000 uses the learned MAC addresses to return packets to those destinations. By default, bridge learning is enabled when you boot up the ASN-9000. To enable bridge learning, issue the following command:

```
learning|learn penable <seglst>
```

**<seglst>** Specifies the segment(s) that you want bridge learning enabled. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here are the results produced by this command:

```
40:ASN-9000:bridge# learning penable 1.4
Learning enabled on segment 1.4
```

To disable bridge learning on the specified segment(s), issue the following command:

```
learning|learn pdisable <seglst>
```

### 11.3.5 Changing the Aging Interval

Use the **aging set** command to set the bridge table aging time. Aging is a mechanism that periodically clears learned entries from the table. Only dynamic entries (entries learned by the software and not configured manually by the user) are aged by the software. Static entries (those created by the user) do not age.

At an interval you specify (the aging interval), the ASN-9000 software determines which of the learned entries in the table have not been recently used. Each learned entry that has not been used during the specified interval is marked aged. This value shows up in the Flags column of the bridge table.

If an entry marked aged is used during the next aging interval, the aged flag is removed and the entry remains in the table. However, if an entry marked aged is unused during the next interval, the entry is removed from the table. Here is the syntax for this command:

```
aging set [<time>]
```

**<time>** Specifies the aging time to clear learned entries in seconds. Default is set to 60 minutes.

To unset aging enter the following command:

```
aging unset
```

### 11.3.6 Deleting an Entry from the Bridge Table

The **bt del** command can also be used to delete a permanent bridge entry from the bridge table. The entry is deleted by issuing this command along with the entry's Ethernet address. Here is the syntax of this command:

```
bt del [<ethaddr>]
```

**<ethaddr>** Specifies the Ethernet address of the entry to be deleted.

## 11.4 Defining and Adding Bridge Groups

Use the **pset group** command to define a network group. A network group is a specific subset of network segments among which packets can be bridged, creating a Layer-2-only VLAN. A packet from one segment in the network group can be bridged only to the other segments in the network group. You can define up to 32 network groups. Group membership can overlap, and each segment can belong to all, some, or none of the network groups.

As shipped from the factory, the ASN-9000 bridging engine contains one network group known as **default**. All segments attached to the ASN-9000 automatically belong to the **default** network group. The **default** group is added to your configuration file when you save the ASN-9000 configuration. If you want to restrict bridging within your network, you can delete the **default** network group and define your own network groups.

#### NOTE

If you save the ASN-9000 configuration using the **system savecfg** command, the default group is automatically added to the configuration file. If your configuration requires that not all segments belong to a common network group (for example, if you defined groups with restricted sets of segments), make sure you delete the default group before saving the configuration file. To delete the default group, issue the following command: **punset group <groupname>**.

Here is the syntax of the **pset group** command:

```
pset group <groupname> <seglist>
```

**<groupname>** Specifies the name of the network group. You can specify any alphanumeric string up to 15 characters in length.

**<seglist>** Specifies the segment(s) that belongs to the network group. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

If you specify **all**, all segments are added to the network group.



To create a new **default** group, you must specify **all** or list all the segments in the ASN-9000 as the <seglist>. If you specify a <seglist>, instead of **all**, and the <seglist> does not include all the segments in the ASN-9000, the software creates a network group called **old\_default**. This default group is stored in the configuration file when you save the configuration.

### 11.4.1 Displaying the Bridge Groups

The **config show** command can be used to display bridge groups configured on the *ForeRunner* ASN-9000. Here is the syntax of this command:

```
config [show] groups
```

Here is an example of configured bridge groups:

```
40:ASN-9000:bridge# config show groups
```

```
Bridge Groups
Name          Segment List
-----
default       1.1, 1.2, 1.3, 1.4, 1.5
               1.6, 2.1, 2.2, 2.3, 2.4
```

## 11.4.2 Deleting a Bridge Group

Use the **punset group** command to delete a bridge group. Here is the syntax for this command:

```
punset group <groupname>

<groupname>    Specifies the name of the network group you want to
                delete.
```

## 11.5 Displaying the Bridging Status of a Segment

Use the **status show** command to display the bridge status for each segment. When you issue this command, the bridge status can differ depending whether you are bridging or routing on particular segments. Here is the syntax for this command:

```
status [show]
```

Here is an example of the display produced by the **status** command:

```
40:ASN-9000:bridge# status show
Segment      Segment Name      Spanning-tree
-----
1.1          Port_1            disabled
1.2**        Port_2            forwarding
1.3          Port_3            disabled
1.4          Port_4            forwarding
1.5          Port_5            blocking
1.6          Port_6            forwarding
2.1          Port_14           forwarding
2.2          Port_15           forwarding
2.3          Port_16           disabled
2.4          Port_17           disabled
```

For bridge or VLAN traffic to be forwarded on the segment, the Spanning-tree state must be forwarding. The Spanning-Tree state does not affect routed traffic on the segment.

Note that the Spanning-tree state blocking does not indicate a problem in your network. As described in Section 11.6.1, the Spanning-Tree algorithm breaks loops in your bridge network by blocking certain segments. The columns in this display show the following information:

<b>Segment</b>	<p>The segment number listed in this column corresponds to the physical location of the segment in the ASN-9000 chassis. Use the <b>system config show</b> command to display information about a segment's physical location in the chassis. See your <i>ForeRunner ASN-9000 Hardware Reference Manual</i>, for more information about this command.</p> <p>If the segment number is followed by ** (two asterisks), then bridging has been disabled by the <b>bridging</b> command on that segment. Note that the bridging command does not affect routing. In this example, bridging has been disabled on segments 1.1, 1.3, 2.3, and 2.4.</p>
<b>Segment Name</b>	<p>The description assigned to each segment. You can change the description using the <b>media sset segment name</b> command. See the <i>ForeRunner ASN-9000 Hardware Reference Manual</i>, for more information about this command.</p>
<b>Spanning-tree</b>	<p>The Spanning-Tree algorithm automatically causes segments to forward or block traffic based on the network topology. When the Spanning-Tree algorithm is enabled, this column shows one of four states:</p> <ul style="list-style-type: none"><li>listening</li><li>learning</li><li>blocking</li><li>forwarding</li><li>disabled</li></ul> <p>The listening and learning states occur when you first enable the Spanning-Tree feature or when your network topology changes. The blocking state indicates that packets are not being forwarded. The forwarding state indicates that packets can be forwarded on the segment. The disabled state indicates that the segment has been disabled using the segment command.</p>

In this example, the Spanning-Tree feature is blocking bridge traffic on segment 1.5. The Spanning-Tree state has no effect on routing. However, this state does affect VLANs because traffic is bridged within VLANs, rather than routed.

## 11.6 Configuring Spanning-Tree Parameters

---

The Spanning-Tree algorithm is a mechanism that logically eliminates physical loops in a bridged network. For example, if your bridges are configured in such a way that broadcast/multicast packets are eventually forwarded back to the bridge that first sent them, your network has a loop. Unless you reconfigure your network topology or your bridges to break the loop, or implement a mechanism to logically break the loop, broadcast/multicast packets are forwarded from bridge to bridge indefinitely, clogging your network. Whenever a segment's state is changed, either by automatic segment-state detection or by a user-interface command, the Spanning-Tree algorithm adjusts the network topology accordingly.

When the Spanning-Tree algorithm is enabled using the `spantree` command (see Section 11.6.1) you can fine-tune the following Spanning-Tree parameters:

- Bridge priority
- Segment priority
- Timer threshold
- Spanning-Tree path cost
- Fast hello-time thresholds (if the fast hello-time feature is enabled)

The first four parameters always are used; the last one is optional. The following sections describe how to adjust these parameters. To display the current settings for these parameters, issue the following command:

```
config [show] st
```

Here is an example of the display produced by the **config [show] st** command:

```
46:ASN-9000:bridge# config show st
Spanning Tree
  Status:           Enabled
  System Priority:   8000
  Spanning Tree Add: 01-80-c2-00-00-00
  My Bridge Address: 00-00-ef-02-42-50
  Max Age:          Curr val: 21  (Config val: 21)
  Hello Time:       Curr val: 4   (Config val: 4)
  Forward Delay:    Curr val: 16  (Config val: 16)
  Send Fast Hellos: Disabled
  Fast Hello Parms: HelloTime:1sec, HighUtil:70%, LowUtil:50%

Seg Prio PathCost DesigBridge DesSeg DesCost StaChngs
--- ---
1.1 128 100      this-bridge 1      10      1
1.2 128 100      this-bridge 2      10      1
1.3 128 100      this-bridge 3      10      1
1.4 128 100      this-bridge 4      10      1
1.5 128 100      this-bridge 5      10      1
1.6 128 100      this-bridge 6      10      1
2.1 128 100      this-bridge 7      10      1
2.2 128 100      this-bridge 8      10      1
2.3 128 100      this-bridge 9      10      1
2.4 128 100      this-bridge 9      10      1
```

### 11.6.1 Enabling or Disabling Spanning Tree

To enable the Spanning-Tree algorithm, issue the following command:

```
spantree enable|disable
```

**enable|disable** Specifies whether you are enabling or disabling the Spanning-Tree algorithm. The default is **disable**.

### 11.6.2 Changing Spanning-Tree Parameters

This section describes how to change the setting of individual Spanning-Tree parameters.

#### 11.6.2.1 Setting the Bridge Priority

Use the **spantree set bridge-priority** command to adjust the bridge priority. Here is the syntax for this command:

```
spantree|st set bridge-priority|bp <priority>
```

**<priority>** Is a hexadecimal number in the range from 0 through **ff**(hex). The default is **80**(hex).

The setting for the bridge priority is displayed in the System Priority field of the **config show st** display.

### 11.6.2.2 Setting a Segment's Priority

Use the **spantree set seg-priority** command to adjust the bridge priority for a segment or a list of segments. Here is the syntax for this command:

```
spantree|st set seg-priority|sp <priority> <seglist>
```

**<priority>** Is a hexadecimal number in the range from 0 through **ff**(hex). The default is **80** (hex). You must specify a separate priority for each segment.

**<seglist>** Specifies the segments for which you are setting the priority. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here is an example of this command. Notice that a separate priority is specified for each segment in the segment range. Even if you plan to assign the same priority to all the segments, you must list the priority individually for each segment.

```
27:ASN-9000:bridge# spantree set seg-priority 0 1.2
Ok. Spanning tree has been reinitialized
```

The segment priority for each segment is displayed in the Priority field of the **config show st** display.

### 11.6.2.3 Setting the Path Cost

You can adjust the Spanning-Tree cost to reconfigure the shortest path to the root bridge. Use the **spantree set path-cost** command to adjust the Spanning-Tree cost on a per-segment basis. Here is the syntax for this command:

```
spantree|st set path-cost|pc <path-cost> <seglist>
```

**<path-cost>** Specifies the cost of the path. You can specify a value from 1 through 65535. The default is 100 for 10Mb/s Ethernet segments, and 10 for Fast Ethernet segments. You must specify a separate path cost for each segment.

**<seglst>** Specifies the segments for which you want to adjust the Spanning-Tree cost. You can specify a single segment or a comma-separated list of segments.

Here is an example of this command. Notice that a separate path cost is specified for each segment in the segment range. Even if you plan to assign the same path cost to all the segments, you must list the path cost individually for each segment. In the following example, the path cost 90 is assigned to segment 1. 3.

```
27:ASN-9000:bridge# set path-cost 90 1.3
```

The path cost for each segment is displayed in the Path Cost field of the **config show st** display.

### 11.6.2.4 Setting the Maximum Age

Use the **spantree set maxage** command to adjust the maximum age of the bridge-timer threshold. Here is the syntax for this command:

```
spantree|st set maxage <time>
```

**<time>** Specifies the maximum age, in seconds. You can specify from 6 through 40 seconds. The default is 21 seconds.

The maximum age of the bridge timer threshold is displayed in the Max Age field of the **config show st** display.

### 11.6.2.5 Setting the Hello Time

Use the **spantree set hello** command to adjust the hello time of the bridge-timer threshold. Here is the syntax for this command:

```
spantree|st set hello <time>
```

**<time>** Specifies the hello time, in seconds. You can specify from 1 through 10 seconds. The default is 4 seconds.

To display the current setting, issue the **config show st** command to display the Spanning-Tree settings, then check the value in the Hello Time field.

### 11.6.2.6 Setting the Forward Delay

Use the **spantree set fwddelay** command to adjust the forward delay of the bridge-timer threshold. Here is the syntax for this command:

```
spantree|st set fwddelay <time>
```

**<time>** Specifies the forward delay, in seconds. You can specify from 4 through 30 seconds. The default is 16 seconds.

The forward delay of the bridge timer threshold is displayed in the Priority field of the **config show st** display.

### 11.6.2.7 Setting the Fast-Hello Time

Under heavy network traffic, Spanning-Tree hello packets are not transmitted at regular hello-time intervals. Such irregular time intervals can delay the transmission of hello packets. If hello packets are delayed past a certain time value, called the maximum age, your Spanning-Tree state can change. If the segment state is “blocking,” and hello packets are not received before another time value, the Max Age, your Spanning-Tree state will change to “listening” and then to “learning.”

Use the **spantree set fast-hello** command to enable or disable the fast hello timer feature. Here is the syntax for this command:

```
spantree|st set fast-hello <time>
```

This feature is disabled by default. To display the current setting, issue the **config show st** command to display the Spanning-Tree settings, then check the value in the Sending Fast Hellos field.

### 11.6.2.8 Setting the High- and Low-Utilization Percentage

If the fast hello timer feature is enabled, when a segment’s utilization exceeds an upper-end value (<high-util>), the *ForeRunner* ASN-9000 software automatically compensates for the increased traffic by using fast hello time to transmit hello packets. The fast hello time is less than the normal (configured) hello time. When all segments’ utilizations drop below a lower-end value, the <low-util>, the hello time reverts to normal (either previously configured or system defaults).

Use the **spantree set high-util** command to set the high-utilization threshold. Here is an example of this command:

```
spantree|st set high-util <percentage>
```

**<percentage>** Specifies the upper-end value of segment utilization. If segment utilization exceeds this value and the fast hello timer feature is enabled, ASN-9000 automatically compensates for the increased network traffic. This value is a percentage in the range of 1 to 100. The default is 70%.

Use the **spantree set low-util** command to set the low-utilization threshold. Here is an example of this command:

```
spantree|st set low-util <percentage>
```

**<percentage>** Specifies the lower-end value of segment utilization. If segment utilization drops below this value, the *ForeRunner* ASN-9000 software automatically reverts to normal hello time. This value is a percentage in the range of 1 to 100. The default is 50%.

## 11.7 Displaying and Clearing Bridge Statistics

---

Use the **stats show** command to display table misses, a subset of statistics, or all statistics. Here is the syntax for this command:

```
stats [show]
```

### 11.7.1 Clearing Statistics

Use the **stats clear** command to clear all bridge statistics. This command resets all bridge statistics to 0, then begins collecting statistics again. Once you clear the bridge statistics, the statistics displayed in response to the **stats** command show the counts since the most recent clear, rather than since the most recent reboot.

## 11.8 Displaying and Clearing the Bridge Cache

---

The *ForeRunner* ASN-9000 software maintains a bridge cache. Each time the bridging engine bridges a packet, it creates an entry in the bridge cache containing the packet's destination Ethernet address and source Ethernet address. The bridge cache is frequently updated with the most-recently used source-destination pairs and provides a fast path for bridge traffic resulting in increased performance.

You can use the bridge cache for at-a-glance information about the current bridge traffic in your network.

## 11.8.1 Displaying the Bridge Cache

To display the bridge cache, issue the following command:

```
cache [show] <seglist>
```

**<seglist>** Specifies segments for which you want to display the cache entries. Specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here is an example of the bridge cache:

```
17:ASN-9000:bridge# display-cache
Bridging cache:
Port 01: Dest: 08-00-20-08-70-54, Source: 08-00-20-0f-dd-99
Dest: 00-00-6b-82-3f-34, Source: 08-00-20-0f-6c-96
Dest: 08-00-20-08-85-69, Source: 08-00-20-0f-dd-99
Dest: 08-00-20-08-70-54, Source: 08-00-20-0f-6c-96
Port 02: Dest: 00-00-6b-82-3f-34, Source: 08-00-20-0e-ae-03
Dest: 00-00-94-06-79-12, Source: 08-00-20-10-56-53
Dest: 08-00-20-0f-f2-9d, Source: 08-00-20-0e-ae-03
Dest: 08-00-20-10-19-ac, Source: 00-00-6b-82-3f-34
Dest: 08-00-20-0f-f2-9d, Source: 00-00-6b-82-3f-34
Dest: 08-00-20-10-06-e3, Source: 00-00-6b-82-3f-34
Dest: 02-cf-1f-90-40-23, Source: 08-00-20-10-56-53
Dest: 08-00-20-08-70-54, Source: 08-00-20-10-56-53
Dest: 00-00-c0-ed-61-4a, Source: 08-00-20-10-56-53
Dest: 08-00-20-08-85-69, Source: 00-00-6b-82-3f-34
Dest: 08-00-20-08-70-54, Source: 08-00-20-0e-ae-03
Listing continues
Port 21: empty
Port 22: empty
Port 23: empty
Port 24: empty
```

## 11.8.2 Clearing the Bridge Cache

To ensure that the entries displayed in the bridge cache are recent and reflect current traffic patterns, you can clear the cache just before displaying it. To clear the bridge cache, issue the following command:

```
cache clear
```

The entire contents of the bridge cache are removed.



# CHAPTER 12

## SNMP Subsystem Commands

The ASN-9000 contains an implementation of Simple Network Management Protocol (SNMP). SNMP uses User Datagram Protocol (UDP), an industry-standard connectionless protocol used to send and receive packets between a managed ASN-9000 and other devices.

This chapter describes the commands in the `snmp` subsystem and tells you how to perform the following tasks:

- Display the SNMP configuration.
- Add an SNMP management community.
- Add an SNMP manager.
- Delete an SNMP management community.
- Delete an SNMP manager.
- Display SNMP packet statistics.
- Delete SNMP packet statistics.

In addition, this chapter describes how to set up files for use with SunNet Manager to access the MIBs.

Using a third-party SNMP application, you can access ASN-9000 MIB (Management Information Base) objects for information about the ASN-9000. The software contains implementation of standard MIBs and the ASN-9000 Proprietary MIB.

### 12.1 Accessing the SNMP Subsystem

---

To access the `snmp` subsystem, issue the following command at the runtime command prompt:

```
snmp
```

## 12.2 Displaying the SNMP Configuration

---

To display the current configuration of communities and managers, use the **config show** command.

```
config [show] [-l] [<community-name>]
```

- l Specifies that managers and trap configurations are to be listed.
- <community-name> Specifies the community for which you want configuration information displayed.

When used with no arguments, this command lists only communities and each community's access (the first and second column from the example shown below).

The example that follows illustrates the **config show** command containing both arguments. It shows the manager and trap configuration for a specific community.

```
47:ASN-9000:snmp# config show -l admin
```

Community	Access	Managers	Traps
-----	-----	-----	-----
admin	rw	147.128.7.3	notrap

## 12.3 Displaying Statistics

---

The SNMP subsystem maintains statistics on SNMP packets that it transmits and receives. These statistics are a superset of the corresponding statistics provided in the SNMP table of MIB-II.

The *ForeRunner* ASN-9000 software maintains two copies of each SNMP statistics counter:

- Count since last clear.
- Count since last switch reset.

To display these statistics, use the following command:

```
stats [show] [-t]
```

- t Displays statistics since last switch reset.

The example that follows shows the **stats** command used without the **[-t]** argument:

```
10:ASN-9000:snmp# stats
SNMP packet statistics (count since last stats clear):
Packets Rcvd:          49086      Packets Sent:          49086
Bad Version Rcvd:      0          Bad Comm Name Rcvd:      0
Bad Comm Uses Rcvd:    0          ASN Parse Err Rcvd:    0
Bad Type Rcvd:         0          Too Big Rcvd:          0
No Such Name Rcvd:     0          Bad Values Rcvd:       0
Read Onlys Rcvd:       0          Gen Errs Rcvd:         0
Total vars Req:        49086      Total vars Set:         0
Get Req Rcvd:          0          GetNext Req Rcvd:      49086
Set Req Rcvd:          0          Get Resp Rcvd:         0
Traps Rcvd:            0          Too Big Sent:          0
No Such Name Sent:     0          Bad Values Sent:       0
Read Onlys Sent:       0          Gen Errs Sent:         0
Get Req Sent:          0          GetNext Req Sent:      0
Set Req Sent:          0          Get Resp Sent:        49806
Traps Sent: 0
```

### 12.3.1 Clearing Statistics

To clear statistics, use the **stats clear** command as in the following example:

```
51:ASN-9000:snmp# stats clear
Okay
```

## 12.4 Adding an SNMP Community

The default switch configuration includes the standard default SNMP community, **public**, which has read-only access. You can add other communities with the **community add** command. Here is the syntax for the **community add** command.

```
community|com add <community-name> [ro|rw]

<community-name>    Specifies the community that you want to add.
[ro|rw]              Specifies the community's access as read-only (ro) or
                     read-write (rw). The default is read-only access.
```

The example that follows illustrates adding an **admin** community with read-write access.

```
43:ASN-9000:snmp# community add admin rw
```

The software supports a maximum of eight SNMP communities.

## 12.4.1 Supported SNMP Traps

The ASN-9000 supports the following standard traps defined in RFC 1157:

- coldStart
- warmStart
- linkDown
- linkUp
- authenticationFailure

## 12.4.2 Deleting an SNMP Community

Communities are deleted using the **community delete** command. Here is the syntax for this command:

```
community|com delete|del <community-name>
```

**<community-name>** Specifies the community that you want to delete.

This command deletes both the community name and all managers that have been associated with it.

The command that follows illustrates deleting the **admin** community.

```
44:ASN-9000:snmp# community delete admin
```

## 12.4.3 Adding an SNMP Manager

Each community can include up to 16 managers. Managers are added with the **manager add** command. Here is the syntax for the **manager add** command:

```
manager|man add <community-name> <IP-addr> [trap|notrap]
```

**<community-name>** Specifies the community name for which you want to add a manager.

**<IP-addr>** Specifies the IP address of the manager.

**trap|notrap** Is an optional flag, indicating whether the manager should receive traps. If the manager should receive traps, use (**trap**). If the manager should not receive traps, use (**notrap**). The default is **notrap**.

This command deletes both the community name and all managers that have been associated with it.

The command that follows illustrates deleting the **admin** community.

```
44:ASN-9000:snmp# community delete admin
```

In the example that follows, a manager with IP address 147.128.7.3 is added to the **admin** community.

```
44:ASN-9000:snmp# manager add admin 147.128.7.3 notrap
```

## 12.4.4 Deleting an SNMP Manager

To delete a manager, use the delete manager command. Here is the syntax for the delete manager command:

```
manager|man delete|del <community-name> <IP-addr>|all
```

- |                               |  |
|-------------------------------|--|
| <b>&lt;community-name&gt;</b> | Is the community name of the manager which you want to delete.   |
| <b>&lt;IP-addr&gt; all</b>    | Is the IP address of the manager which you want to delete. If you use <b>all</b> , all managers in the community are deleted, but the community itself remains configured. |

## 12.4.5 Preparing Files for SunNet Manager

If you plan to use SunNet Manager to access the ASN-9000 MIBs, you must prepare the following types of files for each MIB:

- Schema.
- Trap.
- OID.

Table 12.1 lists the utilities and the file names in SunNet Manager used to prepare these files.

**Table 12.1 - SunNet Manager utilities.**

schema	A MIB converted from ASN.1 format.	mib2schema	<MIB-name>.schema
trap	Active traps for a particular MIB.	mib2schema	<MIB-name>.trap
OID	Object Identify file. Translates the Object Identifiers used by SNMP to communicate into the identifiers that SunNet Manager understands.	mib2schema	<MIB-name>.oid

\*Where <MIB-name> is the name of the MIB.

# CHAPTER 13 IP Subsystem Commands

This chapter describes the commands in the `ip` subsystem and tells you how to use them to configure and manage the ASN-9000 switch as an IP router. Using `ip` subsystem commands, you can:

- Show the switch's IP configuration
- Add, show, and delete IP interfaces
- Enable IP routing (and allocate additional memory to the IP route table)
- Add, show, and delete IP routes
- Enable, show current settings for, and change the configuration of Router-Discovery.
- Show, add static entries to, and delete static entries from the IP Address Resolution Protocol (ARP) table.
- Ping IP workstations or other IP routers
- Add and delete IP helper addresses
- Customize the routing behavior
- Show and clear IP, ICMP, ARP, RIP, and IP Helper statistics
- Show or clear the IP route cache

## 13.1 Accessing the IP Subsystem

---

To access the `ip` subsystem, issue the following command at the runtime command prompt:

```
ip
```

## 13.2 Displaying the IP Configuration

---

You can display the current IP configuration using the **config [show]** command. Here is an example of the display produced by this command:

```
117:ASN-9000:ip# config show
IP Configuration:
-----
IP Forwarding:                enabled (gateway)
Load Balancing:               Off
Default TTL:                  64
Arp cache aging time:         5:00
Routing Network Broadcasts:   enabled
VLAN Bridging Network Broadcasts: enabled
Routing Broadcast Packets:    enabled
Send ICMP redirects:          enabled
Forward Pkts with SrcRt Option: enabled
Arp auto-learn:               enabled (gateway)
Arp VLAN Strict:              enabled (gateway)
Routed Packet Snooping:       disabled
```

You can set any of the IP configuration items listed in this display.

<b>IP Forwarding</b>	Indicates whether IP forwarding is enabled or disabled. See Section 13.3.9.)
<b>Load Balancing</b>	<p>Enables the ForeRunner ASN-9000 switch to distribute IP traffic to remote destinations among up to four equal-cost routes.</p> <p>When load balancing is enabled, up to four load-balancing slots are used per destination to identify next-hop gateways. Packets are hashed to a slot according to source and destination IP address so that packets belonging to a given flow always take the same path.</p>
<b>Default TTL</b>	Indicates the time-to-live (TTL) parameter. This parameter specifies how long a packet is allowed to remain in the net before it is dropped. (See Section 13.8.3.)
<b>ARP cache aging time</b>	Indicates when unused learned entries in the ARP table are removed from the ARP table if they continue to be inactive. (See Section 13.6.2.)
<b>Routing Network Broadcasts</b>	Indicates whether routing of network broadcast packets in a subnetted environment is enabled or disabled. (See Section 13.8.6.2.)

<b>VLAN Bridging Network Broadcasts</b>	Indicates whether bridging of network broadcast packets over a VLAN is enabled or disabled.
<b>Routing Broadcast Packets</b>	Indicates whether routing of network broadcast packets addressed to the ASN-9000 Ethernet address is enabled or disabled. The default is “enabled.”
<b>Send ICMP redirects</b>	Indicates whether ICMP redirect messages are enabled or disabled.
<b>Forward Pkts with SrcRt Option</b>	Indicates whether the software is permitted to forward IP packets containing source route options.
<b>ARP auto-learn</b>	Indicates whether the software is automatically learning ARP entries. (See Section 13.6.)
<b>Routed Packet Snooping</b>	Indicates whether packet snooping is enabled or disabled. Packet snooping lets you examine the IP filtering cache for information about the IP packets routed by the switch.

## 13.3 Configuring and Showing IP Interfaces

---

Before you can use your ASN-9000 switch as an IP router, you must assign an IP address to each segment through which you plan to route IP packets. When discussing TCP/IP, a connection to a physical segment is called an *interface*.

You can assign multiple IP addresses to the same segment. In addition, you can create a VLAN (virtual LAN) by assigning the same IP address to multiple segments. By default, the routing software routes IP packets among different subnets, but bridges IP packets among segments on the same subnet.

When you configure an IP interface (using the **add interface** command), the ASN-9000 software automatically sets the MTU value for the IP interfaces based on the medium type:

- For interfaces on Ethernet segments, the MTU is set to 1500.
- If the interface spans multiple segments, and those segments include Ethernet, the MTU value is set to 1500.

Before you begin configuring your IP interfaces, read the considerations and restrictions in Section 13.3.1 and Section 13.3.2. For information about adding IP interfaces, see Section 13.4.

**NOTE**

If you want to configure the switch to listen to RIP broadcasts on a subnetwork, but you do not want to add an IP interface address to do so, you can add a directly-attached subnet.

### 13.3.1 Considerations

The following considerations apply to assigning interface addresses.

- An interface address must be specified in dotted-decimal notation, and it must be a valid IP host address. A valid IP address must contain a host number that is non-zero and non-broadcast (broadcast IDs are all binary 1s).
- When you add an IP interface, you can specify a subnet mask containing all ones or all zeroes.
- When an interface address is assigned to a segment, the routing software assumes that the segment is physically connected to a net whose IP network number equals the <network-number> part of the interface address. Routing occurs between networks with different network numbers.

**NOTE**

Unlike other switches, the ASN-9000 allows the same IP network number to be assigned to multiple segments (creating a virtual LAN). When this is done, the software bridges IP packets among like-numbered nets that are connected to physically distinct segments.

- The ASN-9000 allows multiple interface addresses with different network numbers to be assigned to a single segment. When this is done, the software forwards packets for any of the corresponding nets to that segment.
- Even if routing is not desired, an interface address must be assigned to a segment in order for TELNET or SNMP connections to be made through that segment. A remote workstation uses this interface address when establishing a TELNET or SNMP connection to the ASN-9000.

### 13.3.2 Restrictions

The following restrictions apply when you assign IP interface addresses. These restrictions are necessary to ensure reliable switch operation. Invalid configurations can bring down an entire network.

- When a single network number appears on multiple segments, all those segments must be assigned the same interface address and subnet mask.
- You cannot configure a parent network address when one or more subnets of that address have been configured on one or more segments.
- You cannot configure a subnet address if its parent network address has been configured on one or more segments. The parent network is the overall network on which subnetworks are configured. For example, network 147.128.0.0 is the parent network of subnetworks 147.128.1.0 and 147.128.2.0. These two subnetworks are referred to as children networks of the parent network.
- You cannot assign different IP host addresses to one interface on the same network or subnet.
- For proper operation under RIP, subnet addresses should normally all have the same binary length—in other words, they should all use the same subnet mask. If you find it necessary to assign variable-length subnet addresses (different subnet masks for some addresses), you must observe certain rules. For information on these rules, refer to the *ForeRunner ASN-9000 Filters Reference Manual*.

### 13.3.3 How the Software Handles IP Packets

- The software discards unexpected IP broadcast packets. The IP broadcast software traps IP broadcast packets and discards them immediately if they were not expected by the switch. This feature is particularly beneficial for large networks that experience high volumes of broadcast traffic.
- The software discards IP broadcast packets that are bridged back to it. Some workstations bridge broadcast packets (including RIP packets) sent from the ASN-9000 back to the switch. The ASN-9000 IP checks the IP source address of the incoming packet to determine whether the packet came from the switch itself. If the packet did come from the switch, the switch discards the packet.
- The software routes IP broadcast packets addressed to the ASN-9000 Ethernet address. If you need to disable routing of IP broadcast packets addressed to the ASN-9000 Ethernet address, use the `enable|disable route-net-broadcast` command. (See Section 13.8.6.)

### 13.3.4 Showing the IP Interface Table

Use the **interface [show]** command to display the IP interface addresses that are configured. For each segment, the table lists the IP addresses assigned to the segment, the link state of the segment (UP or DOWN), and other information. Here is the syntax for this command:

```
interface|it [show] [-s] [<disprestrictors>]
```

**-s** Displays additional statistics, including the number of packets and octets transmitted to and received from the net by each interface.

**<disprestrictors>** Specifies segments for which you want to display the IP addresses. Specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here are some examples of the use of this command.

```
25:ASN-9000:ip# interface show
      Vlan Seg   InterfaceAddr   SubnetMask   bcast   MTU   state   cost
          3   1.3   147.128.132.1   255.255.252.0   br0     1500   down    0
      3   1.4   147.128.132.1   255.255.252.0   br0     1500   up      0
      3   1.5   147.128.132.1   255.255.252.0   br1     1500   down    5
```

By using the **-s** argument with this command, you also can view certain routing statistics, including the number of packets and octets transmitted to and received from the net by each interface.

```
26:ASN-9000:ip# interface show -s
      Vlan Seg   InterfaceAddr   SubnetMask   bcast   MTU   state   cost
          3   1.3   147.128.132.1   255.255.252.0   br0     1500   down    0
0 pkts in, 0 octets in, 0 pkts out, 0 octets out
      3   1.4   147.128.132.1   255.255.252.0   br0     1500   up      0
0 pkts in, 0 octets in, 0 pkts out, 0 octets out
      3   1.5   147.128.132.1   255.255.252.0   br1     1500   down    5
0 pkts in, 0 octets in, 0 pkts out, 0 octets out
```

### 13.3.5 Adding an IP Interface

Use the **interface add** command to assign an IP address to a ASN-9000 segment. When you add an interface address, the software makes an entry into the IP route table to show that the corresponding network is connected to the specified segment. The software then creates the interface. Here is the syntax for this command:

```
interface|it add <vlanid> <ipaddr> [/<prefixlen>|<mask>] [br[oadcast]
0|1] [met[ric] <metric>]
```

<vlanid>	Specifies the VLAN ID you want to assign to the specified segment(s). By assigning the same IP address to multiple segments, you can create a VLAN. The IP address must be in dotted-decimal notation (four decimal numbers in the range 0–255 separated by dots).
<ipaddr>	Specifies the IP address you want to assign to the specified segment(s). The IP address must be in dotted-decimal notation (four decimal numbers in the range 0–255 separated by dots).
<prefixlen>	Allows you to create a valid variable-length subnet by using the <b>interface add</b> command. For more information about variable-length subnets, see Section 13.3.2.
<mask>	Allows a standard IP subnet mask to be used. If a particular network uses IP subnet addressing, then the subnet mask should be specified here using dotted-decimal notation. Otherwise, the switch uses a default subnet mask equal to the “natural” subnet mask for the particular class of address.
[br[oast]0 1]	<p>Specifies the style of broadcast address on a segment-by-segment basis:</p> <p>When you specify <b>br0</b>, the switch sends an “all-0s” broadcast. This means all bits in the host segment of the address are 0s. The <b>br0</b> argument is useful when the ASN-9000 interoperates with workstations that use the old style of IP broadcast address, with all-0s as the host number.</p> <p>When you specify <b>br1</b>, the switch sends a standard “all-1s” broadcast. This means all bits in the host segment of the address are 1s. The default is <b>br1</b>.</p>
[met[ric] <metric>	Specifies an additional cost of using the subnet interface. This cost is the number of extra hops to the destination. The range is 1 through 14. (The router decrements an IP packet’s time-to-live field at each hop.) The default is zero. When the switch reports this subnet using RIP, it adds the additional cost to the reported metric.

The cost parameter can provide controlled routing in the presence of redundant paths, such as when two ASN-9000s are connected in parallel for redundancy. The cost of the attached subnets can be set to a value greater than zero in one of the switch. When you set the cost to a value greater than zero, routing is forced through the other switch if it is alive. See *Chapter 15, IP/RIP Subsystem Commands* for a description of RIP and routing operations with redundant paths.

The example below shows the use of the **interface add** command to add an interface:

```
20:ASN-9000:ip# interface add 1 192.12.20.17
Adding net 192.12.20.0: Okay
Port 1, Addr 192.12.20.17, Mask 255.255.255.0, added
```

Before adding the interface address, the software makes an entry into the route table to show that the corresponding network (192.12.20.0) is directly connected to the specified segment.

The following example shows how to add a single interface address to multiple segments using one command:

```
22:ASN-9000:ip# it add 2.2, 2.4 147.128.132.1 255.255.252.0
Adding subnet 147.128.132.0: Okay
Port 2.2, Addr 147.128.132.1, Mask 255.255.252.0, added
Port 2.4, Addr 147.128.132.1, Mask 255.255.252.0, added
```

You can assign an interface address with a non-zero cost to force routing through a desired path in the presence of redundant paths. In the following example, segments 1 and 2 are physically connected to the same router:

```
22:ASN-9000:ip# it add 1.1 147.128.132.1 255.255.255.0
Adding subnet 147.128.132.0: Okay
Port 1.1, Addr 147.128.132.1, Mask 255.255.255.0 added
23:ASN-9000:ip# it add 1.2 147.128.136.1 255.255.255.0 cost 3
Adding subnet 147.128.136.0: Okay
Port 1.2, Addr 147.128.136.1, Mask 255.255.255.0, cost 3, added
```

Because a higher cost is assigned to segment 1.2, all routing is forced through segment 1.1.



When making changes to IP address or subnet mask, it is not necessary to reboot the ASN-9000.

### 13.3.6 Deleting an IP Interface

Use the `interface del` command to delete one or more interface addresses. Here is the syntax for this command:

```
interface del[ete] [-p] <vlanid>|all <ipaddr>|all
```

- [-p]** Allows you to preserve the address-based parameters of RIP entries.
- <vlanid>|all** Specifies the VLAN IDs for which you want to delete the corresponding interfaces. You can specify a single VLAN ID or a comma-separated list of VLAN IDs. If you specify **all**, all VLANs are deleted from the specified segments.
- <ipaddr>|all** Specifies the IP addresses for which you want to delete the corresponding interfaces. You can specify a single address, a comma-separated list of addresses, or a hyphen-separated range of addresses. If you specify **all**, all IP addresses are deleted from the specified segments.



When you delete the last interface to a particular net, that net is automatically deleted from the route table.

Here is an example of the use of this command. To delete a particular interface address on a particular segment, specify the segment number and interface address.

```
31:ASN-9000:ip# interface delete 3.1 147.128.132.1
Interface address 147.128.132.1, Port 3.1: deleted
```



When making changes to IP address or subnet mask, it is not necessary to reboot the ASN-9000.

### 13.3.7 Configuring VLANs

To make managing your network segments easier, the ASN-9000 lets you create VLANs (Virtual LANs). A VLAN is a network that spans two or more physical segments. VLANs make network configuration changes simple by letting you create and change LANs logically using software commands, as opposed to physically moving segment cables.

You can define any number of segments in the ASN-9000 as members of a VLAN. VLANs can overlap, so the same segments can be members of more than one VLAN. You even can define multiple VLANs on the same segment.

For each segment in the VLAN, the effective bandwidth available to nodes on the VLAN increases. For example, a VLAN containing six 10 Mb/s Ethernet segments enjoys 60 Mb/s of bandwidth. Even though bandwidth is increased, administration and management overhead for the segments in the VLAN does not increase, because the segments can be managed as a single network.

To add a segment to a VLAN, you merely create an interface that contains that segment along with the other segments you want to place in the VLAN. To add a VLAN, issue the **vlan add** command. Here is the syntax for this command:

```
vlan add <vlanid> <seglist>
```

**<vlanid>** Specifies the VLAN IDs for which you want to create the corresponding VLANs. You can specify a single VLAN ID or a comma-separated list of VLAN IDs.

**<seglist>** Specifies the segments for which you want a VLAN. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

In the following example, vlan 2 has been added to segments 1.2, 1.3, and 1.4. The **vlan show** command is used to display the results of the VLAN created:

```
31:ASN-9000:ip# vlan add 2 1.2,1.3,1.4  
32:ASN-9000:ip# vlan show  
2                1.2,1.3,1.4
```

### 13.3.7.1 Changing the VLAN Configuration

Because VLANs are created using software commands, rather than by rearranging your network segments, you can easily change VLANs to suit your needs. To change a VLAN, issue the **vlan tset** command. The following is the syntax for this command:

```
vlan tset <vlanid> seglist <seglist>
```

**<vlanid>** Specifies the VLAN IDs for which you want to change the corresponding VLANs. You can specify a single VLAN ID or a comma-separated list of VLAN IDs.

**seglist** Specifies that you are changing the segments associated with the VLAN.

**<seglist>** Specifies the segments for which you want to change the VLAN. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

or:

```
vlan tset <vlanid> new-name <new vlanid>
```

**<vlanid>** Specifies the VLAN IDs for which you want to change the corresponding VLANs. You can specify a single VLAN ID or a comma-separated list of VLAN IDs.

**new-name** Specifies that you are changing the name of the VLAN.

**<new vlanid>** Specifies the new name you are giving to the corresponding VLAN.

In the following example, segments 1.4 and 1.5 have been added to VLAN 1:

```
31:ASN-9000:ip# vlan show
1                1.2,1.3
32:ASN-9000:ip# vlan tset 1 seglist 1.4-1.5
33:ASN-9000:ip# vlan show
1                1.2,1.3,1.4,1.5
```

In the following example, VLAN 1 has been changed to VLAN 2:

```
31:ASN-9000:ip# vlan show
1                1.2,1.3,1.4,1.5
32:ASN-9000:ip# vlan tset 1 new-name 2
33:ASN-9000:ip# vlan show
2                1.2,1.3,1.4,1.5
```

### 13.3.7.2 Deleting a Configured VLAN

To delete a configured VLAN, issue the `vlan delete` command. Here is the syntax for this command:

```
vlan del[ete] <vlanid>
```

**<vlanid>** Specifies the VLAN IDs for which you want to delete the corresponding VLANs. You can specify a single VLAN ID or a comma-separated list of VLAN IDs.

### 13.3.8 Allocating Memory for Additional IP Routes

Before you begin using the `ip` subsystem, allocate memory for the subsystem by issuing the `addmem` command. The memory allocation increases the capacity of the IP route table. You specify the additional memory in terms of IP routes. The increment is 1K routes. The following example shows the results of this command:

```
1:ASN-9000:ip# addmem
IPR: Routing Table is now : 1 K
2:ASN-9000:ip#
```

If memory has been allocated for IP routing at the time you save the configuration with a `system savecfg` command, the corresponding `ip` subsystem `addmem` command is placed in the configuration file ahead of the other `ip` commands. Thus, you only need to type the `addmem` command when you first configure the ASN-9000 for `ip` routing.

### 13.3.9 Enabling IP Routing

After you define the IP interfaces (see Section 13.3) you are ready to enable IP routing using the following command:

```
enable ip
```

IP routing is disabled by default.

## 13.4 Showing, Adding, and Deleting IP Routes

---

This section describes how to show the IP route table and interpret its contents. This section also describes how to manually add and delete static route-table entries. Note that the software makes additions to the IP route table in two basic ways: it “learns” them from a routing protocol (RIP or OSPF) or the user adds them manually. Learned routes are often called “dynamic entries” and user-added routes are often called “static entries”.

### 13.4.1 Showing the IP Route Table

To display the IP route table, issue the following command:

```
route [show] [-c|-r|-s|-o] [d|t] [-a] [-f] [<seglist>] [<ipaddrlist>]
```

- [-c|-r|-s|-o]** Filters the display according to the type of route:
  - c Displays only directly connected entries.
  - r Displays only RIP routes.
  - s Displays only static routes and directly connected routes.
  - o Displays only OSPF routes.
- [d|t]** Displays additional information, including statistics for packets and bytes. When this argument is specified, the **-f** argument is ignored. **t** displays the total number of routes.
- [-a]** Displays only active routes.
- [-f]** Displays routes that are in the DOWN state.
- <seglist>** Specifies the segments for which you want route information. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.
- <ipaddrlist>** Specifies the IP address list of segments for which you want to display route information. You can use a wildcard character (\*) in place of any part of the IP address.

Here is an example of the use of this command. The **-f** argument is used to display selected routes that are in the DOWN state.

```
35:ASN-9000:ip# route show -f
```

Destination	Subnet Mask	Gateway	Met	State	RtSrc	Age	Port
-----	-----	-----	---	-----	-----	----	----
195.1.1.0	255.255.255.0	-----	0	Down	Direct	----	None

```
Total routes: 1 (Direct: 1, Static Routes: 0, RIP routes: 0)
```

For each IP route, the route table shows the following information:

- Destination**     The IP address of the destination host or net.

<b>Subnet Mask</b>	The subnet mask used by the destination host or net.
<b>Gateway</b>	If the destination is not directly attached to the ASN-9000, this field contains the IP address of the gateway (router) through which packets for the destination are to be routed.
<b>Met</b>	For entries learned through RIP, this field shows how many hops (routers) away the destination is. For example, if a packet must go through one more router to reach its destination, the metric is 1.
<b>State</b>	The state of the route. Possible states are Up or Down, corresponding to active and inactive.
<b>RtSrc</b>	Indicates the source of the routing information: <ul style="list-style-type: none"><li><b>Direct</b> Indicates that the destination is directly attached to the ASN-9000. Such entries are added automatically when you issue the <b>ip interface add</b> command.</li><li><b>OExTp1</b> Indicates that the route was learned through OSPF, from a type-1 External LSA (link-state advertisement). An External LSA indicates that the destination is in another Autonomous System.</li><li><b>OExTp2</b> Indicates that the route was learned through OSPF, from a type-2 External LSA.</li><li><b>OInter</b> Indicates an inter-area route learned through OSPF. The destination is in an area to which the ASN-9000 is connected, but the ASN-9000 is not in the area that contains the destination. OSPF learns the inter-area routes from summary LSAs.</li><li><b>OIntra</b> Indicates an intra-area route learned through OSPF. The destination is in an area that contains the ASN-9000.</li><li><b>Static</b> Indicates that the route was manually added using the ip add-route command.</li><li><b>RIP</b> Indicates that the route was learned through RIP.</li></ul>

<b>Age</b>	Used only by RIP. Indicates how many seconds have passed since fresh information about this route was received.
<b>Port</b>	Lists the segment on which packets for this destination should be forwarded. For directly attached nets, a list of segments can appear, because the ASN-9000 allows a single net to be used on multiple segments.

## 13.4.2 Adding an IP Route

The ASN-9000 stores information about routes in the route table. Entries in the route table are learned dynamically by RIP (as described in *Chapter 15, IP/RIP Subsystem Commands*) or you can configure entries into the table manually (static entries). You can assign static routes for individual hosts or for entire nets.

All nets that have corresponding interface addresses assigned to one or more ASN-9000 segments are considered to be directly attached to the switch. When such interface addresses are assigned by the **interface add** command, the software automatically makes a corresponding entry in the route table. As a result, the routing software automatically routes any incoming IP packet whose destination address is on a directly attached net to the corresponding segment(s). No additional configuration is required.

Additional information is required, however, to route packets to destinations that are not directly attached. In many cases, routers can use RIP to dynamically discover routes that are not directly attached to the hosts and nets. Routes also can be statically assigned, as described in this section. If RIP is not running, routes to non-directly-attached hosts and nets must be assigned statically. To assign the route to be used when forwarding to a host or net, use the **route add** command.

```
route add [-s] [-d] <destination> <gw-ipaddr> <metric> <segment>
```

<b>[-s]</b>	Specifies the addition of a strict route.
<b>[-d]</b>	Specifies the addition of a route in the DOWN state.
<b>&lt;destination&gt;</b>	The IP address of the destination host or net.
<b>&lt;gw-ipaddr&gt;</b>	Specifies the IP address of the gateway (router) to which packets destined for the specified host are forwarded. Generally, this gateway is connected to the ASN-9000 through a net. The net is directly attached to both the gateway and the switch.

<b>&lt;metric&gt;</b>	The cost of the route (number of hops to the destination). Generally, the route used is the one with the lowest cost, regardless of whether it is static (added to the route table permanently by the <b>route add</b> command) or learned through RIP.
<b>&lt;segment&gt;</b>	Specifies the ASN-9000 segment onto which a packet is forwarded to reach the specified gateway and the host.

Here is an example of how this command is used.

```
61:ASN-9000:ip#rt add 192.9.208.1 255.0.0.0 147.128.128.65 3 5
192.9.200.1 255.0.0.0, 147.128.128.65, 3, 5: Added-route is active
```

### 13.4.3 Enabling and Disabling Load Balancing

When load balancing is enabled, the ASN-9000, receiving packets from the same source, uses different routes for the incoming packets to reach the ASN-9000 without any delay. To enable load balancing on the ASN-9000, issue the following command:

```
load-balance|lb enable|disable
```

<b>enable disable</b>	Specifies whether you are enabling or disabling load balancing. The default is disable.
-----------------------	---

### 13.4.4 Enabling Loopback Detection

When loopback detection is enabled, the ASN-9000 sends a special loopback-detect packet on each outbound segment that has at least one IP address. To enable loop detection on the ASN-9000, issue the following command:

```
loop-detection|ld enable|disable
```

<b>enable disable</b>	Specifies whether you are enabling or disabling loop detection.
-----------------------	---

Here is an example of this command:

```
69:ASN-9000:ip# ld enabled
loop-detection: enabled
```

#### 13.4.4.1 Setting the Loopback Detection Time

To set the loopback detection time, issue the following command:

```
loop-detection|ld set time <value>
```

**<value>** Specifies the time interval in minutes for sending out loopback-detection packets. The default is 10 minutes.

Here is an example of this command:

```
70:ASN-9000:ip# ld set time 15
71:ASN-9000:ip#
```

13.4.4.2 Displaying the IP Loop Detection Table

To display the IP loop detection table, issue the following command:

```
loop-detection|ld [show]
```

Here is an example of this command:

```
72:ASN-9000:ip# ld show
loop-detection:                               enabled
  IP Loop Detection Table:
IP Address      MAC Address      TTL      rport      Segment(s)
-----
147.128.128.2   08-00-20-08-70-54   16       2          1.3
```

For each IP route, the route table shows the following information:

- IP Address** The IP address of the outbound segment sending the loopback-detect packet.
- MAC Address** The Ethernet address of the host.
- TTL** Specifies how long a packet is allowed to remain in the net before it is dropped. Packets that cannot find or are blocked from their destination nodes are dropped when the TTL expires.
- rport** Specifies the receiving port of the segments sending the loopback-detect packet.
- Segment (s)** The segment(s) that sent the loopback-detect packet.

13.4.5 Enabling or Disabling an IP Route

After you define the IP interfaces (see Section 13.3) you are ready to enable IP routing using the following command:

```
route enable|disable <destination> <gw-ipaddr> <metric> <segment>
```

**enable|disable** Specifies whether you are enabling or disabling IP routing. The default is disable.

<b>&lt;destination&gt;</b>	The IP address of the destination host or net.
<b>&lt;gw-ipaddr&gt;</b>	Specifies the IP address of the gateway (router) to which packets destined for the specified host are forwarded. Generally, this gateway is connected to the ASN-9000 through a net. The net is directly attached to both the gateway and the switch.
<b>&lt;metric&gt;</b>	For entries learned through RIP, this field shows how many hops (routers) away the destination is. For example, if a packet must go through one more router to reach its destination, the metric is 1.
<b>&lt;segment&gt;</b>	Specifies the segments for which you want to enable or disable the feature. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

## 13.4.6 Deleting an IP Route

You can delete a static route using the **route del** command. To delete a remote, host-specific entry, issue this command:

```
route del <destination> <gw-ipaddr>
```

<b>&lt;destination&gt;</b>	The IP address of the destination host or net.
<b>&lt;gw-ipaddr&gt;</b>	Specifies the IP address of the gateway.

You cannot use the **route del** command to delete learned entries from the route table. The software automatically removes learned entries that remain unused for 180 seconds.

## 13.5 IP Router Discovery

---

Based on Internet Control Message Protocol (ICMP) to enable hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers, Router Discovery allows hosts to discover routers automatically through a series of solicitation and advertisement messages. This eliminates the need for the specific configuration of static addresses.

Before a host can send IP datagrams beyond its directly-attached subnet, it must discover the address of at least one operational router on that subnet. Typically, this is accomplished by reading a list of one or more router addresses from a (possibly remote) configuration file at start-up time. On multicast links, some hosts also discover router addresses by listening to

routing protocol traffic. Routing Discovery on the ASN-9000 system uses a pair of ICMP [10] messages, for use on multicast links. More information about Router Discovery can be found in RFC 1256. To enable Router Discovery on your ASN-9000, issue the following command:

```
rdm nenable <ipaddr>
```

**<ipaddr>** Specifies the IP address of the host.

To enable Router Discovery on your ASN-9000, issue the following command:

```
rdm ndisable <ipaddr>
```

### 13.5.1 Setting the Advertisement Address

You can specify whether to send out advertise messages to the all-systems multicast address, 224.0.0.1, or to the limited-broadcast address, 255.255.255.255. By default, the ASN-9000 designates the all-systems multicast address. To set the Advertisement Address for Router Discovery on your ASN-9000, issue the following command:

```
rdm nset AdvertisementAddress multicast|bro <ipaddr>
```

<b>AdvertisementAddress</b>	Specifies the IP destination address to be used for multicast Router Advertisements sent from the interface. The only permissible values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255.
<b>multicast bro</b>	Specifies advertise messages to all-systems multicast address, 224.0.0.1, or to the limited-broadcast address, 255.255.255.255. The ASN-9000 default is all-systems multicast, 224.0.0.1.
<b>&lt;ipaddr&gt;</b>	Specifies the IP address belonging to the interface from which this message is sent, or 0.

## 13.5.2 Setting the Advertisement Preference

A Router Advertisement includes a preference level for each advertised router address. When a host must choose a default router address that is, when, for a particular destination, the host has not been redirected or configured to use a specific router address, the host is expected to choose from those router addresses that have the highest preference level. To set the advertisement preference for Router Discovery on your ASN-9000, issue the following command:

```
rdm nset preference <preference> <ipaddr>
```

<b>preference</b>	Specifies you are setting the preference value for Router Discovery.
<b>&lt;preference&gt;</b>	<p>Specifies the preference of each Router Address as a default router address, relative to other router addresses on the same subnet. A signed, twos-complement value; higher values mean more preferable.</p> <p>A 32-bit, signed, twos-complement integer, with higher values meaning more preferable. The minimum value (hex 80000000) is used to indicate that the IP address, even though it may be advertised, is not to be used by neighboring hosts as a default router address.</p>
<b>&lt;ipaddr&gt;</b>	Specifies the IP address belonging to the interface from which this message is sent, or 0.

## 13.5.3 Setting the Advertisement Interval

A Router Advertisement also includes a lifetime field, specifying the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements. This is used to ensure that hosts eventually forget about routers that fail, become unreachable, or stop acting as routers. The default advertising rate is once every 10 minutes, and the default lifetime is 30 minutes. To set the advertisement interval for Router Discovery on the ASN-9000, issue the following command:

```
rdm nset interval <time> <ipaddr>
```

<b>interval</b>	Specifies you are setting the interval value for Router Discovery.
<b>&lt;time&gt;</b>	Specifies the time allowed between sending multicast Router Advertisements from the interface.

**<ipaddr>** Specifies the IP address belonging to the interface from which this message is sent, or 0.

### 13.5.4 Displaying the Advertisement Interval

To display the Router Discovery table, issue the following command:

```
rdm [show]
```

Here is an example of the results produced by this command:

```
72:ASN-9000:ip# rdm show
-- RDM Configuration --
IP Address      ADVERTISEMENT      Interval      Preference      this
                Address                                     Level          IP address
-----
200.1.1.1       multicast            10:00         0               yes
150.1.1.2       multicast            10:00         0               yes
```

## 13.6 Showing and Configuring the ARP Table

The ASN-9000 IP routing software maintains an ARP table of IP-to-Ethernet address translations. These translations are used to route packets and, under some circumstances, to generate replies to ARP requests. There are three ways that entries are added to the ARP table:

- When a host uses ARP to request the ASN-9000 Ethernet address, the host’s IP and Ethernet addresses are recorded (“learned”).
- If a host forwards a packet to a destination through the ASN-9000, it can generate an ARP request to learn the destination’s Ethernet address. When the switch receives the reply to such a request, it records the destination’s IP and Ethernet addresses.
- Permanent entries are added using ASN-9000 commands.

### 13.6.1 Enabling and Disabling ARP

To enable ARP, issue the following command:

```
arp enable auto-learn
```

**auto-learn** Indicates enabling auto-learn of incoming packets on the ASN-9000. Default is auto-learn enabled.

To disable ARP auto-learning, issue the following command:

```
arp disable auto-learn
```

### 13.6.2 The ARP Cache

The ASN-9000 IP software queues IP route packets for which the ARP table does not contain entries, then sends an ARP request to learn the Ethernet address of the destination device. When the ARP reply is received from the destination device, the queued packet is forwarded. The source node does not need to send the packet to the switch again.

### 13.6.3 Showing the ARP Table

Use the **arp [show]** command to display the current contents of the ARP table. Here is the syntax for this command:

```
arp [show] [-r] [-t] [-s] [<disp-restrictors>]
```

**[-r]** Specifies raw entries with hash indices and displacements.

**[-t]** Specifies that only the total count of entries is to be displayed.

**[-s]** Specifies that the ARP entries to be displayed are sorted by the IP address (in increasing order).

**<disp-restrictors>** Specifies segments for which you want to display the IP addresses. Specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here are some examples of the use of this command. If no argument is given, then the entire table is displayed, for example:

```
70:ASN-9000:ip# arp show
IP Addr      Ethernet Address  Flags      Segment
147.128.128.2 08-00-20-08-70-54 perm publish    6
147.128.128.3 08-00-20-08-85-69                2
192.9.201.1   02-cf-1f-90-40-23                12
192.9.201.7   08-00-20-0f-dd-99 perm                10
```

Permanent entries can have a switch flag, indicating that the entry was added automatically by ASN-9000, and a broadcast flag, indicating that the Ethernet address is broadcast or multi-cast.

You can specify an optional IP address with the **arp [show]** command, in which case only the entry for that address is displayed:

```
70:ASN-9000:ip# arp show 147.128.128.2
IP Addr      Ethernet Address  Flags      Segment
147.128.128.2  08-00-20-08-70-54      6
```

A “wildcard” character (\*) can be used in place of any byte(s) of the IP address, in which case only entries that match that address are displayed.

### 13.6.4 Clearing the ARP Table

Use the **arp clear** command to clear the ARP table. All learned entries are removed, but static entries (created using the **arp add** command) remain in the table. These must be removed manually using the **arp del** command.

You can use this command to help restabilize the network after a host is moved from one segment to another. When there is activity on the network, the cleared entries quickly reappear in the ARP table, and a host that has been moved will be relearned on its new segment.

### 13.6.5 Showing and Changing the ARP Aging Interval

By default, the ASN-9000 automatically checks learned entries in the ARP table every five minutes to see if they have been used. Each unused entry is marked aged. If an aged entry is used during the next five-minute interval, the aged flag is removed. However, aged entries that remain unused during the second five-minute interval are removed from the ARP table. You can change the aging interval or turn off aging using the **arp set age** command. Here is the syntax for this command:

```
arp set|show|unset age <time>

<time>    Specifies (in minutes) a new aging interval or turns
           aging off. The default is 5 minutes. You must set
           aging time at a minimum of 1 minute (enter either 60
           (seconds) or 1:00). To specify minutes, you must
           specify <minutes:seconds>.
```



If you turn ARP aging off, the ARP table can quickly overflow. Make sure you monitor the table frequently if you turn ARP aging off.

Here is an example of the use of this command:

```
73:ASN-9000:ip# arp set age 30:00
ARP cache aging set to 30 minutes
```

To display the current ARP aging interval, issue the **config show** command.

### 13.6.6 Adding a Static Entry to the ARP Table

Use the **arp add** command to add a static ARP entry to the ARP table. Static ARP entries are not subject to aging and are not cleared when the ARP table is cleared (using the **arp clear** command).

```
arp add [-p] <ipaddr> <ethaddr> <seglist>
```

- |           |  |
|-----------|--|
| [-p]      | If this argument is present, then the <i>ForeRunner</i> ASN-9000 IP routing software replies directly to ARP requests for this entry. Note that this facility is provided only for permanent, not learned, entries in the ARP table.   |
| <ipaddr>  | Specifies the IP address to be translated.   |
| <ethaddr> | Specifies the Ethernet address corresponding to the given IP address.  |
| <seglist> | Specifies the segments to which packets sent to the IP address specified by <ethaddr> are forwarded. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments. In addition, you can specify <b>all</b> to add the ARP entry for all segments. |

Some examples of using the **arp add** command are shown below:

```
75:ASN-9000:ip# arp add -p 147.128.128.8 08-00-02-03-04-05 1.4
Added/changed:
  IP address: 147.128.128.8
  Ethernet address: 08:00:02:03:04:05
  Flags: PERMANENT PUBLISH
Segments: none
```

Permanent and published entries are flagged in the ARP table:

```
77:ASN-9000:ip# at
IP Addr      Ethernet Address  Flags           Segments
147.128.128.1 08:00:02:03:04:05 perm publish    1.1
```

### 13.6.7 Deleting a Static Entry from the ARP Table

Use the **arp delete** command to delete static or dynamically learned ARP entries from the ARP table. Here is the syntax for this command:

```
arp delete <ipaddr>

<ipaddr>    Specifies the IP address in the ARP entry you want to delete.
```

When a host is moved from one segment to another, you can use this command to delete its obsolete learned entry from the ARP table without disturbing any other entries. The network can then relearn the host's new location without being forced to relearn other host locations, as it would if you cleared the ARP table.

## 13.7 Pinging Other IP Devices

The ASN-9000 supports the echo facility of the Internet Control Message Protocol (ICMP) in two ways:

- The ASN-9000 generates a response to any ICMP echo request packet received on any segment.
- Under your control, the ASN-9000 can send an ICMP echo request packet to any IP address.

ICMP echo requests are commonly used to determine whether devices are reachable on the network. UNIX workstations provide a **ping** command that generates an ICMP echo request to a specified IP address.

When you issue the **ping** command from your workstation, the ASN-9000 responds and you can determine whether the ASN-9000 is reachable from your workstation. However, depending upon the configuration, the switch might be known by multiple IP addresses. Unless the workstation is directly connected to the switch, the IP address specified in the **ping** command can affect the route taken, and therefore the reachability of the switch.

Similarly, the ASN-9000 itself provides a ping command to generate an ICMP echo request to a specified IP address. Here is the syntax for this command:

```
ping|pi [-t <timeout>] [-size <size>] <ipaddr>
```

- |                             |  |
|-----------------------------|--|
| <b>[-t &lt;timeout&gt;]</b> | Specifies how many seconds the <i>ForeRunner</i> ASN-9000 switch waits for a response from the specified device. The default is 5 seconds. |
| <b>[-size &lt;size&gt;]</b> | Specifies the packet length. You can specify any length from 64 through 1472 bytes. Default is 64.   |
| <b>&lt;ipaddr&gt;</b>       | Specifies the IP address of the device you are trying to reach.  |

Here are some examples of the use of this command.

```
83:ASN-9000:ip# ping 147.128.128.8
147.128.128.8 is alive
84:ASN-9000:ip# ping 147.128.128.15
No response from 147.128.128.15
```

The **ping** command normally waits 5 seconds for the specified host to respond before timing out. However, you can specify a shorter or longer time-out, as shown in the following example. In this example, a one-second delay is specified.

```
85:ASN-9000:ip# ping 147.128.128.8 1
No response from 147.128.128.8
```

## 13.8 IP Helper

---

This section describes how to use the IP Helper feature. IP Helper is an enhancement to the **ip** subsystem that assists client stations on one network segment in communicating with servers on another network segment when the two segments are connected by a ASN-9000. This includes situations where one switch, as a client station, needs to boot from a server from which it is separated by another switch. By default, the IP Helper feature is configured to help packets destined for any of the following standard UDP ports:

- BootP client packets (port 68).
- BootP server packets (port 67).
- Domain Name System (port 53).
- IEN-116 Name Server (port 42).
- NetBIOS Datagram Server (port 138).
- NetBIOS Name Server (port 137).
- TACACS service (port 98).
- TFTP (port 69).
- Time service (port 37).

If you need to add a UDP port to this list, you can do so using the **helper add -d** command. (See Section 13.8.2.)

## 13.8.1 How IP Helper Works

When a client sends out a broadcast packet addressed to a server that is directly connected to the client, the server:

- Receives the limited broadcast IP packet sent out by the client.
- Uses the client's Ethernet address to look up its corresponding IP address.
- Sends a unicast packet in reply.

This also is true if the client and server are on different segments, but the segments are defined as part of the same virtual LAN. In this case, the packets are bridged.

However, if the client and server are on different segments separated by a router (gateway), the client's broadcast packet never reaches the server. If the intervening router is an ASN-9000, you can use the IP Helper facility on that ASN-9000 to tell it where to forward UDP packets sent by the client.

To use IP Helper to help a client reach its server, assign the server's IP address as an IP Helper address to the ASN-9000 segment connected to the client. When this segment receives a UDP packet from the client, it forwards the packet to the node that has the IP address corresponding to the ASN-9000 segment's IP Helper address.

For the UDP packet to be successfully forwarded, the following criteria must be met:

- The packet must be received on a segment where an IP Helper address is configured.
- The destination UDP port must be in the UDP-helper Port Table on the router.

See RFC 1542 for more information.



IP Helper does not affect the forwarding of limited-broadcast packets in a virtual LAN environment. The same packet can be forwarded to multiple segments that are on the same virtual LAN.

## 13.8.2 Using IP Helper

Before you can use IP Helper:

- The ASN-9000 must be configured as an IP router. (See Section 13.3.9.)
- An IP Helper address must be assigned to the segment which connects to the diskless workstation or other device that is being helped. The IP Helper address is the address of the desired server on the network.

To display helper configuration on an IP segment, issue the following command:

```
config [show] helper
```

### 13.8.2.1 Adding an IP Helper Address

To add an IP Helper address to a segment, issue the following command:

```
helper add <IPaddr> [<UDPportlist>] <seglist>  
helper add -d
```

<IPaddr>	Specifies the helper address. You must specify the IP address of the server as the helper address.
[ <UDPportlist>]	Specifies any of the standard UDP ports available by default.
<seglist>	Specifies the segments on which you want to add an IP address. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments. If the segment list is set to <b>all</b> , then the IP address is assigned to all valid segments in the switch.
-d	Displays the contents of the default UDP portlist. Allows you to specify additional default UDP ports.

Here is an example of the use of this command.

```
11:ASN-9000:ip# helper add 147.128.42.37 1.4  
Helper address 147.128.42.37: added.
```

An IP Helper address is added to segment 1.4 on the router switch. When used, this IP Helper address routes UDP packets received on segment 1.4 to IP address 147.128.42.37.

You can assign multiple IP Helper addresses to a single segment, or multiple segments to a single IP Helper address. Assigning multiple IP Helper addresses to a single segment provides redundancy when multiple servers are used.

### 13.8.2.2 Deleting an IP Helper Address

To delete an IP Helper address, issue the **helper delete** command. Here is the syntax for this command:

<b>helper delete &lt;IPaddr&gt; &lt;UDPportlist&gt; default[s] all &lt;seglist&gt;</b>	
<b>&lt;IPaddr&gt;</b>	Specifies the helper address. You must specify the IP address of the server as the helper address.
<b>&lt;UDPportlist&gt; default[s] all</b>	Specifies the type of UDP port you are deleting. If you specify <b>all</b> , all UDP ports, including the default ports are deleted.
<b>&lt;seglist&gt;</b>	Specifies the segment(s) which connect the router to the client ASN-9000. If you specify <b>all</b> , all entries assigned to the specified IP address are deleted.

Here is an example of the use of this command.

```
16:ASN-9000:ip# helper delete 2.2.2.2 1.2
2.2.2.2:138 (netbios-dgm), port 1.2 :deleted
2.2.2.2:138 (netbios-ns), port 1.2 :deleted
2.2.2.2:138 (tacnews), port 1.2 :deleted
2.2.2.2:138 (tftp), port 1.2 :deleted
2.2.2.2:138 (bootps), port 1.2 :deleted
2.2.2.2:138 (dns), port 1.2 :deleted
2.2.2.2:138 (name), port 1.2 :deleted
2.2.2.2:138 (time), port 1.2 :deleted
```

### 13.8.2.3 Displaying Statistics and the UDP Table

To display current statistics for an IP Helper address defined for a segment, issue the **show-helper** command. A table is displayed listing the segment, helper address, the number of packets helped, and the number of packets dropped. Here is the syntax for this command:

<b>helper show [-p -s]</b>	
<b>helper show -d</b>	
<b>[-p -s]</b>	Allows you to sort the IP Helper table by UDP port <b>-p</b> , or by segment number <b>-s</b> .
<b>-d</b>	Displays the contents of the default UDP portlist. Allows you to specify additional default UDP ports.

Here is an example of the use of this command.

```
11:ASN-9000:ip# helper show
Helper IP      UDP port Segment   Helped   Reverse   Dropped
-----
147.128.48.37  37 time    1.4        0         0         0
```

The table in this example shows that during the current session, IP Helper address 147.128.48.37 has helped four UDP packets (perhaps BOOTP packets) find their IP destinations. The table also shows that one UDP packet was dropped. Note that the **helper show** command lists statistics only for those UDP packets that the switch tried to help.

UDP packets can be dropped for any of the following reasons:

- The helping ASN-9000 switch does not have a route to the destination address in the UDP packet.
- The helping ASN-9000 switch runs out of resources to redirect the packet.

In addition, for BOOTP packets only, the following conditions can cause the helping *ForeRunner* ASN-9000 switch to drop the packet:

- The hop count in the packet has been exceeded.
- A gateway has already helped the packet. (A bit in the packet is set when the packet is helped.)

### 13.8.2.4 Deleting Default UDP Entries

To delete default UDP entries, issue the **helper delete** command. Here is the syntax for this command:

```
helper delete -d <UDP ports to remove>
```

- d Displays the contents of the default UDP portlist. Allows you to delete additional default UDP ports.

### 13.8.2.5 Clearing Statistics

To clear the IP Helper statistics, issue the **stats clear helper** command. Here is an example of the use of this command.

```
12:ASN-9000:ip# stats clear helper
IP helper table stats are cleared.
```

### 13.8.2.6 Deleting an IP Helper Address

To delete an IP Helper address, issue the del-helper command. Here is the syntax for this command:

```
helper delete <IPaddr> [<UDPportlist>|default[s]|all <seglist>
helper delete -d <UDP ports to remove>
```

<b>&lt;IPaddr&gt;</b>	Specifies the helper address. You must specify the IP address of the server as the helper address.
<b>[&lt;UDPportlist&gt; default[s] all</b>	Displays the contents of the default UDP portlist. If the UDP port list is set to <b>all</b> , then the IP address is deleted from all valid segments in the switch.
<b>&lt;seglist&gt;</b>	Specifies the segment(s) which connect the router to the client switch. If you specify <b>all</b> , all entries assigned to the specified IP address are deleted.
<b>-d</b>	Displays the contents of the default UDP portlist. Allows you to specify additional default UDP ports to delete.

If both *<seglist>* and *<IPaddr>* are specified as **all**, all IP Helper definitions on the router switch are deleted.

### 13.8.2.7 Adding an IP Helper Gateway IP Address

To add an IP Helper gateway IP address to a segment, issue the following command:

```
helper add -g <GwIP-Add> <seglist>
```

<b>&lt;GWIP-Add&gt;</b>	Specifies the defined IP address to be used as the gateway address while helping a bootp/DHCP packet. A maximum of 10 gateway addresses can be configured per segment. When more than one gateway address is configured for a segment, all the gateway addresses will be used sequentially.
<b>&lt;seglist&gt;</b>	Specifies the segments on which you want to add a gateway IP address. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.
<b>-g</b>	Specifies gateway address configuration of the available gateway IP addresses.

Here is an example of the use of this command.

```
11:ASN-9000:ip# helper add -g 147.128.42.37 1.4
Helper address 147.128.42.37: added.
```

A gateway IP Helper address is added to segment 1.4 on the router switch. When used, this IP address will be used as the gateway address when a packet being helped is routed.

You can assign multiple IP Helper gateway addresses to a single segment, or multiple segments to a single IP Helper gateway address.

### 13.8.2.8 Deleting an IP Helper Gateway Address

To delete an IP Helper address, issue the **helper delete** command. Here is the syntax for this command:

```
helper delete -g <GwIP-Add> <seglist>
```

- |                         |   |
|-------------------------|---|
| <b>&lt;GwIP-Add&gt;</b> | Specifies the defined IP address used as the gateway address to help a bootp/DHCP packet you want to delete.  |
| <b>&lt;seglist&gt;</b>  | Specifies the segments on which you want to delete a gateway IP address. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments. |
| <b>-g</b>               | The contents of the available gateway IP addresses to delete.   |

### 13.8.2.9 Displaying IP Helper Gateway Addresses

To display an IP Helper address, issue the **helper show** command. Here is the syntax for this command:

```
helper show -g
```

- |           |   |
|-----------|---|
| <b>-g</b> | Displays the contents of the available gateway IP addresses to display. |
|-----------|---|

Here is an example of the use of this command.

```
11:ASN-9000:ip# helper show
Segment  Configured GW_Addrs
-----  -
1.4      147.128.42.37
```

### 13.8.3 Setting the Time-To-Live Parameter

The time-to-live (TTL) parameter specifies how long a packet is allowed to remain in the net before it is dropped. Packets that cannot find or are blocked from their destination nodes are dropped when the TTL expires. To change the default TTL, issue the following command:

```
ipdefaultttl|ittl set <value>
```

**<value>** Specifies the new TTL time in hops. Specify a number between 1 and 255. The default is 16 hops.

To display the TTL value for outgoing IP packets, issue the following command:

```
ipdefaultttl|ittl [show]
```

### 13.8.4 Enabling and Disabling ICMP Redirect Messages

Use the **send-icmp-redirect** command to enable or disable sending ICMP redirect messages by the ASN-9000.

In networks that use multiple routers, ICMP redirect messages in from routers of alternative routes to segments connected to the routers. Normally, this feature helps optimize routing throughput by ensuring that routers are informed of the most efficient paths to the segments on the network.

The ASN-9000 works well when it receives ICMP redirect messages; however, some other switches do not work well in environments in which these messages are used. If your network contains switches that do not work well when they receive ICMP redirect messages, you can disable sending these messages on the ASN-9000.

```
enable|disable send-icmp-redirect|sir
```

**enable|disable** Specifies whether you are enabling or disabling ICMP redirect messages. The default is **en1** (enabled).

### 13.8.5 Enabling or Disabling Source-Route Filtering

Use the **fwd-pkts-with-srcrt-options** command to disable the source-route feature and strengthen the “firewall” protecting your network from outside users.

IP packets that contain the loose-source-route or the strict-source-route option are forwarded by default in software version 3.0. The source-route options are intended to help forwarding of IP packets. When a packet containing a source-route option is forwarded, the packet can appear to receiving devices as though it originated from the device that forwarded it. As a result, these devices are more likely to accept the forwarded packets, rather than filter them.

Disabling the source-route feature prevents outside users from using and exploiting the source-route contained in packets to gain access to your network. The syntax for this command is:

```
enable|disable fwd-pkts-with-srcrt-option|fps
```

**enable|disable** Specifies whether you are enabling or disabling source-route filtering. The default is **en1** (enabled).



For additional information on IP filtering, see the *ForeRunner ASN-9000 Filters Reference Manual*.

### 13.8.6 Enabling or Disabling Network-Broadcast Forwarding

By default, the ASN-9000 forwards broadcast packets onto subnets attached to the ASN-9000. A network broadcast packet is a packet containing either all zeros or all ones in the host portion of the address. For example: 1.120.255.255, 192.9.200.0, and 10.255.255.255 all are network broadcast packets. The way the software handles broadcast packets differs depending upon how they are received and the destination address specified in the packets.

You can cause the ASN-9000 to forward or drop IP network-broadcast packets sent to subnetted interfaces, by enabling or disabling bridge-net-broadcast and route-net-broadcast:

- The bridge-net-broadcast state affects network-broadcast packets received in Ethernet-broadcast packets. If bridge-net-broadcast is enabled, these packets are forwarded. If bridge-net-broadcast is disabled, these packets are dropped.
- The route-net-broadcast state affects network-broadcast packets received in Ethernet-unicast packets. If route-net-bcast is enabled, these packets are forwarded. If route-net-bcast is disabled, these packets are dropped.

The bridge-net-bcast and route-net-bcast states are completely independent of each other. You can enable both or one only, or disable both, depending upon the level of broadcast traffic you want to allow for subnetted interfaces.

IP network-broadcast and IP subnet-broadcast packets can be encapsulated in one of the following types of packets:

- Ethernet-broadcast packets. These packets contain (encapsulate) IP subnet-broadcast packets or IP network-broadcast packets. Ethernet-broadcast packets contain the Ethernet broadcast address (ff-ff-ff-ff-ff-ff) in the destination address field and are received by the ASN-9000 from a directly-attached node.
- Ethernet-unicast packets. These packets contain the ASN-9000s Ethernet address in the destination field. Like Ethernet-broadcast packets, Ethernet-unicast packets can contain (encapsulate) IP subnet-broadcast packets or IP network-broadcast packets. However, unlike Ethernet-broadcast packets, Ethernet-unicast packets are received by the ASN-9000 from another router.

You can selectively enable or disable forwarding of the following types of IP network-broadcasts:

- IP network-broadcasts sent from a node directly-attached to the switch and addressed to a subnetted interface configured on the switch. If bridge-net-bcast is enabled, the packets are bridged to all segments belonging to all subnets in the destination network. If bridge-net-bcast is disabled, the packets are dropped.
- IP network-broadcasts sent from another router and addressed to a subnetted interface configured on the switch. If route-net-bcast is enabled, the packets are routed to all segments belonging to all subnets in the destination network. If route-net-bcast is disabled, the packets are dropped.

Neither the bridge-net-bcast state nor the route-net-bcast state has any effect on IP subnet-broadcast packets or broadcast packets sent to interfaces that are not subnetted:

- If the interface is subnetted and the received packet is a subnet-broadcast, the packet is unconditionally bridged to all the segments belonging to the same subnet.
- If the interface is not subnetted and the received packet is a network broadcast, the packet is unconditionally bridged to all the segments belonging to the same network.
- If the interface is subnetted, and the received packet is a subnet-broadcast, the packet is unconditionally forwarded (routed) to all the segments in the subnet.
- If the interface is not subnetted, and the received packet is a net-broadcast packet, the packet is unconditionally forwarded (routed) to all segments in the network.

### **13.8.6.1 Disabling Bridging of Net Broadcasts**

To prevent the software from forwarding IP network-broadcast packets from directly-attached nodes to subnetted interfaces on the ASN-9000, issue the following command:

```
disable bridge-net-broadcast|bnb
```

After you issue this command, network-broadcast packets encapsulated in Ethernet-broadcast packets are still received internally by the ASN-9000, if applicable, but dropped without being forwarded to the destination subnets. Network-broadcast packets received in Ethernet-unicast packets are not affected.

To re-enable the software to forward network-broadcast packets received in Ethernet-broadcast packets and addressed to subnetted interfaces, issue the following command:

```
enable bridge-net-broadcast|bnb
```

### **13.8.6.2 Disabling Routing of Net Broadcasts**

To prevent the software from forwarding IP network-broadcast packets from other routers to subnetted interfaces attached to the ASN-9000, issue the following command:

```
disable route-net-broadcast|bnb
```

After you issue this command, network-broadcast packets encapsulated in Ethernet-unicast packets are still received internally by the switch, if applicable, but dropped without being forwarded to the destination subnets. Network-broadcast packets received in Ethernet-broadcast packets are not affected.

To re-enable the software to forward network-broadcast packets received in Ethernet-unicast packets and addressed to subnetted interfaces, issue the following command:

```
enable route-net-broadcast|bnb
```

## **13.8.7 Enabling Proxy ARP**

The ASN-9000 supports proxy ARP (RFC 1027), a well-defined mechanism in the TCP/IP protocol suite. Using proxy ARP, a router can respond to an ARP request with its own Ethernet address if it knows a route (or default route) to the destination network or subnet on which the requested address resides.

Without proxy ARP, the requesting host needs to have knowledge of its own network, as well as the destination network and the subnet mask, so that it can ARP the destination directly if it is on the same net or ARP the ASN-9000 (or other gateway) if the destination is on a different net.

Use the proxy-arp command to enable or disable the proxy ARP feature for all segments or a specific list of segments. Here is the syntax for this command:

```
proxy-arp penable|pdisable [<seglist>]
```

**<seglist>** Specifies the segments for which you want to enable or disable the feature. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

**penable|pdisable** Specifies whether you are enabling or disabling the feature. The default is disable (for all segments).

If you do not specify a *<seglist>* or **penable|pdisable**, the current status (enabled or disabled) of the proxy ARP feature is shown.

### 13.8.7.1 Displaying the Proxy ARP Table

To display the Proxy ARP table, issue the following command:

```
proxy-arp [show]
```

Here is an example of the results produced by this command:

```
71:ASN-9000:ip# proxy-arp show
```

```
Segment 2.1: disabled
Segment 2.2: enabled
Segment 2.3: enabled
Segment 2.4: enabled
Segment 2.5: disabled
Segment 2.6: enabled
Segment 2.7: disabled
Segment 2.8: disabled
Segment 2.9: disabled
```

## 13.9 Showing and Clearing Statistics

---

The `ip` subsystem maintains statistics on ARP, ICMP, and general IP packets. These statistics are a superset of the corresponding statistics provided in the SNMP MIB. Use the `stats [show]` command to display statistics. Here is the syntax for this command:

```
stats [show] [-t] [arp|icmp|ip|helper]
```

**-t** Specifies to display all statistics collected since the software was rebooted, rather than just the statistics collected since the last time the `stats clear` command was issued.

**[arp|icmp|ip|helper]** Specifies the type of packet protocol for which you want to display statistics.

The ASN-9000 maintains two copies of each IP statistics counter (and similarly for ICMP and ARP packets):

- Count since last clear.
- Count since last switch reset.

Both counters are updated when the corresponding events occur, but the `stats clear` command clears only the count since last clear. To display the count since last reset, use the `-t` option with the `stats` command.

Here are some examples of the information displayed by the `stats` command. Notice that the first line in each example informs you that statistics since the last statistics clear are being displayed, rather than total statistics accumulated since the switch was last rebooted.

```
IP statistics: count since last stats clear
```

```
Number of Cache Flushes: 1
```

As shown in this example, the IP statistics are organized according to incoming packets and outgoing packets. In addition to totals for packets received, sent, and forwarded, the `stats ip` display lists statistics for many of the types of IP routing errors that can occur in a network.

In the following example, ARP statistics are displayed.

```
74:ASN-9000:ip# stats arp
ARP statistics: count since last stats clear
ARP Packet Statistics:
  Requests received:      38
  Replies received:      25
  Invalid opcodes received: 0
  Requests sent:         226
  Replies sent:          36 (0 proxies)
```

Here is an example of the ICMP statistics displayed by the **stats** command.

```
74:ASN-9000:ip# stats icmp
ICMP statistics: Count Since last stats clear
Messages received:          0   Errors received:          0
Dest unreachable msgs rcvd: 0   TTL expired msgs rcvd: 0
Param prob msgs rcvd:       0   Src quench msgs rcvd: 0
Redirect msgs rcvd:         0   Echo request msgs rcvd: 0
Echo reply msgs rcvd:       0   Timestamp req msgs rcvd: 0
Timestamp repl msgs rcvd:    0   AddrMask req msgs rcvd: 0
AddrMask repl msgs rcvd:     0
Messages sent:              200  Errors sent:          0
Dest unreachable msgs sent: 200  TTL expired msgs sent: 0
Param prob msgs sent:       0   Src quench msgs sent: 0
Redirect msgs sent:         0   Echo request msgs sent: 0
Echo reply msgs sent:       0   Timestamp req msgs sent: 0
Timestamp repl msgs sent:    0   AddrMask req msgs sent: 0
AddrMask repl msgs sent:     0
```

### 13.9.1 Clearing Statistics

Use the **stats clear** command to clear statistics. Here is the syntax for this command:

```
stats clear [arp|icmp|ip|helper|all]
```

**[arp|icmp|ip|helper]** Specifies the type of packet protocol for which you want to clear statistics.

## 13.10 Showing or Clearing the IP Route Cache

The IP routing software maintains a route cache containing translation information for the destination hosts. This information is frequently updated based upon incoming packets on each segment. You can use the route cache to determine which hosts are most frequently used.

### 13.10.1 Displaying the Route Cache

Use the **cache show** command to display the route cache. Here is the syntax for this command:

```
cache show [<seglist>]
```

**[<seglist>]** Specifies the segments for which you want to display the route cache. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

## **13.10.2 Flushing the Route Cache**

You can flush (clear) the route cache using the **cache clear** command. The **cache clear** command removes all entries from the route cache for some or all segments.

After the cache is flushed, new entries are added using the cache's usual most-recently-used algorithm. If you issue a subsequent **cache show** command, fresh entries are displayed.

# CHAPTER 14 IP Multicast Subsystem Commands

This chapter describes the IP Multicast commands and how to use them to define IP interfaces on a ASN-9000 as end stations for IP Multicasting. Unlike IP broadcasting, which sends packets to all destinations, or IP unicasting, which lets you send packets to a single destination, IP Multicasting lets you address and deliver packets to a specific subset of destinations.

**NOTE**

The ASN-9000 IP Multicasting software does not support mrouterd 3.8.

Using the `ip/mcast` subsystem, you can set up and use IP Multicasting with video conferencing and other multicast applications. Using `ip/mcast` commands, you can:

- Show the switch's IP Multicast configuration
- Add, show, and delete IP Multicast interfaces
- Add and delete IP Multicast tunnels
- Enable IP Multicast routing
- Show and clear the IP Multicast route table
- Show and clear the IP Multicast route cache
- Show and clear IP Multicast statistics
- Enable multicast-aware bridging (for systems that perform IP Multicast routing on VLANs)

## 14.1 Accessing the IP Multicast Subsystem

---

To access the `ip/mcast` subsystem, issue the following command at the runtime command prompt:

```
ip/mcast
```

### 14.1.1 Allocating Memory

Before you begin using the `dec` subsystem, allocate memory for the subsystem by issuing the `getmem` command, as shown in the following example:

```
1:ASN-9000:ip/mcast# getmem
Memory allocated for IP Multicast.
2:ASN-9000:ip/mcast#
```

If memory has been allocated for IP Multicast at the time you save the configuration with a `system savecfg` command or `tftp savecfg` command, the corresponding `getmem` command is placed in the configuration file ahead of the other IP Multicast configuration commands. Thus, you only need to type the `getmem` command when you first configure the *ForeRunner* ASN-9000 switch for IP Multicast routing.



FORE Systems recommends that you allocate memory for the IP Multicast subsystem immediately after you boot the ASN-9000 to ensure that the memory you request is available. For more information, refer to the *ForeRunner ASN-9000 Hardware Reference Manual*.

You cannot deallocate memory. To free allocated memory, make sure the configuration file does not contain a `getmem` command, then reboot the software.

### 14.1.2 Enabling Pruning

To enable or disable pruning in the IP Multicast subsystem, issue the following command:

```
pruning enable|disable
```

Here are the results of this command:

```
311:ASN-9000:ip/mcast# pruning enable
IP Multicast pruning enabled.
312:ASN-9000:ip/mcast#
```

## 14.2 Showing the IP Multicast Configuration

You can display the current IP Multicast configuration by issuing the `config show` command. Here is an example of the information shown by the `config show` command.

```
44:ASN-9000:ip/mcast# show config
  IP Multicast Forwarding: disabled
  Multicast Aware Bridging in a VLAN: disabled
  IPM Pruning: enabled
  Max Routing Entries allocated: 2k

Port State for Multicast Traffic:
Segment  2.1:  Disabled ***
Segment  2.2:  Enabled
Segment  2.3:  Enabled
Segment  2.4:  Enabled
```

In this example, the display produced by the `show config` command shows:

- IP Multicast forwarding is enabled.
- Multicast Aware Bridging in a VLAN is disabled.
- IPM pruning is enabled.
- Maximum routing entries allocated is 2k.
- IP Multicast traffic is enabled on all segments except 2.1.

### 14.2.1 IP Considerations

IP Multicast routing works whether IP forwarding is enabled or disabled. In this respect, the ASN-9000 implementation is similar to `mrouted`, which allows multicast routing on a UNIX workstation even if it is not routing regular IP traffic.

#### NOTE

You must enable IP Multicast routing even if the ASN-9000 is configured to have the same subnet on all the segments. The IP Multicast routing code bridges packets intelligently based on reception of membership reports. IP Multicast traffic is restricted to those networks that have listening hosts.

The virtual interface table used for IP Multicast routing is associated closely with the IP interface table. When a virtual interface is added, appropriate information is automatically copied from the IP interface table.

The ASN-9000 also updates the segment list in the virtual interface table whenever you add or delete a segment in an IP interface entry. When you delete an IP interface entry, the software automatically deletes all the IP Multicast virtual interfaces that match the deleted entry's address.

### 14.2.2 Displaying IP Multicast Groups

Use the **multicast-groups show** command to list the IP Multicast group addresses currently known to the ASN-9000 (local router). Here is an example of the display produced by this command.

```
35:ASN-9000:ip/mcast# multicast-groups show
Virtual I\F- : Locaddr: 147.128.70.30 RmtAddr :----, type: Physical
Groups: 224.2.138.32 Segs: 2.1
```

```
Virtual I\F- Locaddr: --- RmtAddr:147.128.90.33, type: Tunnel
```

This table contains the list of IP Multicast groups for each virtual interface. This table contains the following information:

<b>Locaddr</b>	Displays additional statistics, including the number of packets and octets transmitted to and received from the net by each interface.
<b>RmtAddr</b>	Lists the IP address of a remotely attached IP Multicast neighbor. This applies only to tunnels, in which the ASN-9000 and the other end of the virtual interface are separated by gateways.
<b>type</b>	Lists the type of IP Multicast interface. Valid types are Physical and Tunnel.
<b>Groups</b>	Lists the IP Multicast groups. For each group are listed the group IP address and the ASN-9000 segment(s) on which membership reports for that group were received.

### 14.2.3 Displaying IP Multicast Neighbors

Use the **neighbors show** command to list all the neighboring routers currently known to the managed switch. Here is an example of the display produced by this command.

```
35:ASN-9000:ip/mcast# neighbors show
Virtual I\F- :Locaddr: 147.128.128.99 RmtAddr :----,type:Physical,
Neighbors: 147.128.128.30 (25 sec) 147.128.100.2 (40 sec)
Virtual I\F- Locaddr: 147.128.128.99 Rmtaddr---,type:Tunnel,
Neighbors: 130.1.5.1 (35 sec)
```

This display contains the list of neighboring routers for each virtual interface. This table contains the following information:

<b>Locaddr</b>	Lists the IP address of a directly-attached IP Multicast neighbor. This applies only to physical interfaces, in which the ASN-9000 and the other end of virtual interface are directly attached.
<b>RmtAddr</b>	Lists the IP address of a remotely attached IP Multicast neighbor. This applies only to tunnels, in which the ASN-9000 and the other end of the virtual interface are separated by gateways.
<b>type</b>	Lists the type of IP Multicast interface. Valid types are <code>Physical</code> and <code>Tunnel</code> .
<b>Neighbors</b>	Lists the IP Multicast neighbors. For each neighbor are listed the router's IP address and the number of seconds elapsed since the last routing update was received from this neighbor.

## 14.3 Configuring and Showing IP Multicast Interfaces

A physical interface allows two directly connected ASN-9000s (local routers) to communicate with each other. To define a physical interface, use the **interface add** command. The syntax for the command is:

```
it|interface add <ipaddr> [met[ric]<metric>] [thresh[old]<thresh>]
```

<b>&lt;ipaddr&gt;</b>	Is an IP address on the local switch, written in dotted-decimal notation. The address must be present in the IP interface table.
<b>[met[ric]&lt;metric&gt;]</b>	Specifies an additional cost (measured in hops to the destination) of using the interface. The cost range is from 1 through 31. The default is 1.
<b>[thresh[old]&lt;thresh&gt;]</b>	Specifies the minimum time-to-live (TTL) value that an IP Multicast packet must have before it is forwarded over this interface.  This parameter lets you restrict the types of IP Multicast traffic that go out on a network. The default is 1.

Here is an example of the use of the **interface add** command:

```
32:ASN-9000:ip/mcast# interface add 192.10.30.33
Okay
33:ASN-9000:ip/mcast#
```

### 14.3.1 Displaying the Interface Table

Use the **interface [show]** command to display a list of configured virtual interfaces. The display includes both physical interfaces and tunnels. Here is the syntax for this command:

```
it|interface [show] [<disprestrictors>]
```

**[<disprestrictors>]** Specifies segments for which you want to display the IP addresses. Specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

Here is an example of the interface table displayed by the **interface [show]** command.

```
33:ASN-9000:ip/mcast# it
IP Multicast Routing: virtual interface table:
LocalAddress RemoteAddress Type SrcRt Metrc Thrsh State Segments
-----
147.128.70.30 ----- Phy ----- 1 1 Up 2.4,1.8
147.128.128.99----- Phy ----- 1 1 Up 2.3
147.128.33.5 130.1.5.1 Tunl No 1 6 Up 2.4,1.8
147.128.30.30 ----- Phy ----- 1 1 Up 1.5
147.128.33.5 192.9.200.21 Tunl Yes 1 6 Up 1.6
147.128.33.5 ----- Phy ----- 1 1 Up 1.2
```

The table in the example on the previous page contains information about four physical interfaces and two tunnels. The tunnel to destination 130.1.5.1 is an encapsulation tunnel. The tunnel to destination 192.9.200.21 is a source-route tunnel.

<b>Local Address and Remote Address</b>	Identifies the two ends of a tunnel. The local address corresponds to the configured address for a physical interface.
<b>Type</b>	Identifies whether the virtual interface is either a tunnel or a physical interface.
<b>SrcRt</b>	Identifies the type of tunnel. Yes in this column indicates that the tunnel is a source-route tunnel. No in this column indicates that the tunnel is an encapsulation tunnel.
<b>Metrc</b>	Lists the cost (in hops) of the interface.
<b>Thrsh</b>	Lists the threshold value for the interface.

<b>State</b>	Indicates the state of the interface. Up indicates the interface is active. Down indicates the interface is inactive. The interface is DOWN when you disable a segment from the bridging subsystem, or if disabled by the automatic segment-state detection mechanism. See your <i>ForeRunner ASN-9000 Hardware Reference Manual</i> for further information on automatic segment-state detection.
<b>Ports</b>	Lists the segments to which the listed virtual interface is assigned.

### 14.3.2 Deleting a Physical Interface

Use the `interface del` command to delete a physical interface.



When you delete a physical interface, any corresponding tunnels are not deleted. To delete a tunnel, you must use the `tunnel del` command.

Here is the syntax for the `interface delete` command:

```
it|interface delete <ipaddr>|all
```

**<ipaddr>|all** Specifies the IP address of the physical interface to be deleted.

If you specify **all**, all physical interfaces (excluding the tunnels) are deleted.

### 14.3.3 Deleting a Tunnel

Use the `tunnel del` command to delete a virtual interface that maps to a tunnel. Here is the syntax for this command:

```
tunnel del (loc[al]<local-addr> rem[ote]<remote-addr>)|all
```

**loc[al]<localaddr>** Specifies the IP address of the *ForeRunner* ASN-9000 switch (the local end of the tunnel).

**rem[ote]<remoteaddr>** Specifies the IP address of the router at the remote end of the tunnel.

To delete all IP Multicast tunnels, issue the following command:

```
tunnel del all
```

## 14.4 Configuring and Showing Tunnels

---

A tunnel is a type of virtual interface that allows the ASN-9000 (local router) to communicate with a remotely attached router.

### 14.4.1 Adding a Tunnel

Use the **tunnel add** command to define a tunnel. Here is the syntax for this command:

```
tunnel add [-s] loc[al]<local-addr> rem[ote]<remote-addr>  
           [met[ric]<mv>] [thresh[old]<tv>]
```

<b>[-s]</b>	Specifies that the tunnel is a source-route tunnel, rather than an encapsulation tunnel.  If you do not specify <b>-s</b> , this command automatically configures your tunnel as an encapsulation tunnel.
<b>loc[al]&lt;localaddr&gt;</b>	Specifies the IP address of the local router ( <i>ForeRunner</i> ASN-9000 switch). The address must be one of the configured IP addresses listed in the ASN-9000 IP interface table.
<b>rem[ote]&lt;remoteaddr&gt;</b>	Specifies the IP address of the router at the other end of the tunnel.
<b>[met[ric]&lt;mv&gt;]</b>	Specifies an additional cost (extra hops to the destination) of using the virtual interface with which this tunnel is associated. Specify a number in the range 1 through 31. The default is 1.
<b>[thresh[old]&lt;tv&gt;]</b>	Specifies the minimum time-to-live (TTL) value that an IP Multicast packet must have before it can be forwarded through the tunnel. This parameter restricts IP Multicast datagrams from going out on a network. The default is 1.

Here is an example of how to add a tunnel. In this example, the **-s** argument is not used, so the software creates an encapsulation tunnel. The default values are accepted for the metric and threshold.

```
34:ASN-9000:ip/mcast#tunnel add loc 192.10.30.33 rem 155.10.23.222 met 3 thresh 4
Okay
35:ASN-9000:ip/mcast#
```

## 14.4.2 Deleting a Tunnel

Use the **tunnel del** command to delete a virtual interface that maps to a tunnel. Here is the syntax for this command:

```
tunnel del (loc[al]<local-addr> rem[ote]<remote-addr>)|all
```

<b>loc[al]&lt;localaddr&gt;</b>	Specifies the IP address of the ASN-9000 (the local end of the tunnel).
<b>rem[ote]&lt;remoteaddr&gt;</b>	Specifies the IP address of the router at the remote end of the tunnel.

To delete all IP Multicast tunnels, issue the following command:

```
tunnel del all
```

## 14.5 Enabling IP Multicast Routing

To enable IP Multicast forwarding, use the **enable ipm** command. The syntax for this command is:

```
enable|disable ipm
```

<b>enable disable</b>	Specifies whether you are enabling or disabling IP Multicast forwarding. The default is <b>dis</b> .
-----------------------	--



If IP Multicast forwarding is enabled immediately after enabling RIP listening, the multicast route updates are not accepted over a tunnel until the IP routing table learns either an entry to the remote end of the tunnel or a default route.

## 14.5.1 Enabling Multicast Traffic on a Segment

*ForeRunner* ASN-9000 software lets you restrict IP Multicast forwarding on a segment-by-segment basis. The syntax for the command used to enable or disable IP Multicast traffic on a set of segments is shown below.

```
penable|pdisable transmit <segment-list>
```

<b>penable pdisable</b>	Specifies whether you are enabling or disabling IP Multicast forwarding. The default is <b>penable</b> .
<b>&lt;segment-list&gt;</b>	Specifies the list of segments on which IP Multicast forwarding is being enabled or disabled. If you specify <b>all</b> , IP Multicast forwarding is enabled or disabled on all segments.

The first command in the example that follows uses the **transmit** command to enable IP Multicast traffic on segments 2.4. The second command in the example uses the **set** command to disable IP Multicast traffic on segment 2.2.

```
46:ASN-9000:ip/mcast# penable transmit 2.4
Ok
47:ASN-9000:ip/mcast# pdisable transmit 2.2
Ok
```

## 14.6 Configuring and Showing IP Multicast Routes

Use the **route show** command to display a list of IP addresses originating IP Multicast traffic, currently known by the IP Multicast routing software. Here is the syntax for this command:

```
route|rt [show] [-c|-r] [-d] [-t] [-s] [<seglist>] [<ipaddr>]
```

<b>[-c -r]</b>	<b>-c</b> displays directly connected routes only. <b>-r</b> displays DVMRP routes only.
<b>[-d]</b>	Displays the routing table in detail.
<b>[-t]</b>	Displays the total number of routes only.
<b>[-s]</b>	Displays the output in sorted order.
<b>&lt;seglist&gt;</b>	Specifies the ASN-9000 segments for which you want to display route information.
<b>&lt;ipaddr&gt;</b>	Specifies the IP address (origin) of the route entries to be displayed.

Here is an example of the display produced by the command.

```
52:ASN-9000:ip/mcast# route show
IP Multicast Routing table:
Origin          Origin Mask      Gateway          Met  Age  Par.  Seg/Child
-----
147.128.70.0    255.255.255.0    -----
147.128.128.0   255.255.255.0    -----
147.128.90.0    255.255.255.0    -----
129.155.80.0    255.255.240.0    147.128.70.2    3    20   4     2.5,6.2
150.233.0.0     255.255.0.0      147.128.128.111 5    35   2     4.5,2.6
```

The route table contains the following information:

<b>Origin</b>	Lists the IP address of the origin network. An origin is a network that is capable of originating IP Multicast traffic.
<b>Origin Mask</b>	Lists the origin mask used on the origin network. An origin mask is the subnet mask of an origin network.
<b>Gateway</b>	Lists the IP address of the next-hop router to the origin. This column is not applicable to directly connected entries.
<b>Met</b>	Displays the total cost (or metric) of reaching the origin. This metric is the sum of the cost of the next-hop virtual interface and the number of hops or intervening routers (if applicable) used to reach the origin.
<b>Age</b>	Shows the time elapsed, in seconds, since a DVMRP route report was last received for this origin. This column is not for directly connected routes.
<b>Parent</b>	Shows the segment on which the next-hop router is located for a dynamic route. This column is not applicable to directly connected routes.  For directly-connected routes, the Seg/Children column shows the segments on which the corresponding virtual interface is configured. For a dynamic entry, this column lists the segments on which the IP Multicast packets from this origin are forwarded.

## 14.6.1 Clearing the Route Table

Use the **route clear** command to flush all dynamically learned entries from the route table. Here is an example of the use of this command.

```
66:ASN-9000:ip/mcast# route clear
Okay
```

## 14.7 Using the IP Route Cache

---

The IP Multicasting software maintains a route cache containing translation information for the destination hosts. This information is frequently updated based upon incoming packets on each segment. You can use the route cache to determine which hosts are most frequently used.

Because the contents of the route cache can change rapidly, successive **cache show** commands can give different results.

### 14.7.1 Displaying and Clearing the Route Cache

Use the **cache show** command to display the route cache. Here is the syntax for this command:

```
cache show
```

You can flush (clear) the route cache using the **cache clear** command. This command removes all entries from the route cache for some or all segments.

After the cache is flushed, new entries are added using the cache's usual most-recently-used algorithm. If you issue a subsequent **cache show** command, fresh entries are displayed.

## 14.8 Displaying Statistics

---

The `ip/mcast` subsystem maintains statistics on DVMRP, IGMP and routed packets. To display statistics, issue the following command:

```
stats [show] [-t] [dvm|igmp|rt|all]
```

**[-t]** Displays statistics collected subsequent to the last system reset, rather than merely the last time statistics were cleared.

<b>[dvm igmp rt all]</b>	Displays the type of packet for which you want statistics:
<b>dvm</b>	Displays DVMRP packet statistics.
<b>igmp</b>	Displays IGMP packet statistics.
<b>rt</b>	Displays routing packet statistics.

Here is an example of the display produced by the **stats show dvm** command, used to display DVMRP statistics.

```
5:ASN-9000:ip/mcast# stats show dvm
DVMRP Statistics (count since last stats clear):
Route reports sent:                                32
Neighbor probes sent:                               0
Neighbor prunes sent:                               0
Neighbor graphs sent:                               0
Neighbor graft_acks sent:                           0
Neighbor responses sent:                            12
Neighbor2 responses sent:                           0

Route reports received:                             211
Neighbor Probes received:                           1
Neighbor prunes received:                           0
Neighbor graphs received:                           0
Neighbor graft_acks received:                        0
Neighbor requests received:                          0
Neighbor2 requests received:                        33

Rcvd pkts with bad metric:                          1
Rcvd pkts with bad orig. mask:                      0
Rcvd conflicting route reports:                     0
Rcvd truncated route reports:                       0
Conflicting routes deleted:                         0
Rcvd reports from non neighbor:                     5
Rcvd probes from non neighbor                      5
Rcvd prunes from non neighbor                      5
Rcvd grafts from non neighbor                      5
Rcvd graft_acks from non neighbor                   0
Rcvd invalid neighbor requests:                     0
Rcvd invalid neighbor responses:                    0
Rcvd invalid Neighbor2 responses:                   0
Rcvd message from non neighbor:                     0

No mem to receive packet:                           2
No memory to send packets:                           0
```

Here is an example of the display produced by the **stats show igmp** command, used to display IGMP statistics.

```
60:ASN-9000:ip/mcast# stats show igmp
IGMP Statistics (count since last stats clear):
total packets received:                551
short packets received:                2
pkts rcvd with checksum error:         0
total membership queries rcvd:         12
invalid membership queries rcvd:        0
total membership reports rcvd:         333
invalid membership reports rcvd:        0
rcvd packets too big:                  0
rcvd unknown DVMRP message:            0
rcvd unknown IGMP message:            0
packets looped back:                   9
no buffer for looping back:            0
no timers for IP Multicast routing:     0
report not sent - no interface:        0
group timer not started - no I/F:      0
rcvd report from non adj. host:        1
total membership queries sent:         9
total packets sent:                    159
total packets not sent:                 0
no memory to process rcvd pkts:        2
Queue blocks accessed:                  2
Queue blocks released:                  2
Free Queue blocks available:            2048
```

Here is an example of the display produced by the **stats show rt** command, used to display routing statistics.

```
59:ASN-9000:ip/mcast# stats show rt
Multicast routing statistics (count since last clear):
route cache hits:                                661
route cache misses:                             661
route cache flushed:                             0
route lookups:                                   661
route cache misses:                             661
source group pair cache lookups:                11322
source group pair cache misses:                11322
rcvd msg over invalid tunnel:                    5
no room for tunnel options:                      0
rcvd msg on wrong interface:                     17
packets forwarded:                               3213
packets dropped:                                2448
packets received:                               5661
rcvd packet format error:                       0
encapsulated packets rcvd:                      2112
rcvd port not configured:                       0
no route to origin:                             2448
packets bridged:                                1123
packets not bridged:                            0
no memory to process packets:                   0
```

### 14.8.1 Clearing Statistics

Use the **stats clear** command to clear statistics for DVMRP, IGMP, or route packets. Here is the syntax for this command:

```
stats clear [dvm|igmp|rt|all]
```

## 14.9 Enabling Multicast-Aware Bridging

The ASN-9000 supports VLANs (Virtual LANs) for IP routing. A VLAN is an IP interface configured on multiple segments.

When the ASN-9000 receives a packet on an IP Multicast virtual interface that maps to multiple physical segments, it can bridge the packet to other segments and simultaneously route it to other virtual interfaces, transmitting the same copy of the packet on all segments. When this occurs, the time-to-live (TTL) of the bridged packets, as well as the routed packets, is reduced by one. Because this procedure avoids copying the packet again, it results in improved performance. Because most multicast applications use a large TTL value, reducing by one hop when bridging occurs should not significantly affect performance.

## *IP Multicast Subsystem Commands*

If you do not want the ASN-9000 IP Multicasting software to make its forwarding decisions upon the receipt of membership reports on a port, you can enable Multicast bridging by issuing the following command:

```
enable multicast-aware-bridging
```

To disable Multicast bridging, issue the following command:

```
disable multicast-aware-bridging
```

The default is **disabled**.

# CHAPTER 15 IP/RIP Subsystem Commands

In a routed environment, routers communicate with each other to keep track of available routes. The ASN-9000 routing software implements standard RIP (Routing Information Protocol) for exchanging TCP/IP route information with other routers.

RIP uses User Datagram Protocol (UDP), an industry-standard connectionless protocol, for sending and receiving packets between the ASN-9000 and other devices.

This chapter describes how to use `ip/rip` subsystem commands to perform the following tasks:

- Display the switch's RIP configuration
- Configure RIP parameters for IP networks
- Display and clear RIP statistics
- Enable RIP Bridging (used only when IP traffic crosses the ASN-9000 on a VLAN)

## 15.1 Accessing the RIP Subsystem

To access the `ip/rip` subsystem, issue the following command at the runtime command prompt:

```
ip/rip
```

## 15.2 Displaying the RIP Configuration

Use the `config show` command to display the current RIP configuration for a specified IP interface address. Here are the results from this command:

```
57:ASN-9000:ip/rip# config show
RIP Configuration
-----
```

```
RIP Bridging: enabled
```

I/F	Addr	TX	RX	Poison	Rpt Static	Rpt Def	Acpt Def	Auth Type	Key ID	Txtype	Rxtype
19.0.0.1	yes	yes	yes	yes	yes	yes	yes	none	---	rip2	both

## 15.2.1 Configuring RIP Parameters

You can use the **rip nenable** command to set one or more requested RIP parameters for an IP interface address and add the information to the RIP update control table. The syntax for this command is:

```
rip nenable <params> <ipaddr>
```

<b>&lt;params&gt;</b>	Specifies either a comma-separated list of parameters or <b>all</b> for all parameters. Parameters are:
talk   ta	Specifies that the ASN-9000 adds an entry to the RIP packets for the specified subnet.
listen   li	Specifies that the ASN-9000 listens to RIP packets received on the specified IP interface.
poison   po	<p>When a learned route from the specified IP interface goes down, specifies one of the following actions:</p> <p>If <b>no</b> is specified, the ASN-9000 software stops reporting the route.</p> <p>If you specify <b>yes</b>, the software reports route one more time, but with a metric (hop count) of 16, which is infinity as far as RIP is concerned. Other routers learn immediately that the route is down.</p>
RptStaticRt   rs	Specifies whether static routes for the specified IP interface address are reported in RIP packets sent out the segment containing the interface.
RptDefaultRt   rd	When RIP packets are generated on the interface, specifies whether the ASN-9000 reports the default route, if any, in its route table.

AccptDefaultRt | ad Specifies whether default routes for the specified IP interface address are reported to the RIP update control table.


**NOTE**

If the *<params>* is not entered, all parameters except **poison** are set to **yes**.

Here is an example of the use of the **rip nenable** command:

```
56:ASN-9000:ip/rip#rip nenable ta,li,po,rd,rs 19.0.0.1
Okay
57:ASN-9000:ip/rip#
```

If you have already set RIP parameters and decide to change the RIP control type from per-VLAN to per-segment, the ASN-9000 provides a warning message before allowing the change. The change of control type automatically clears the per-VLAN RIP parameters from all pertinent tables. The new parameters are supplied from one of two sources:

- If you saved your configuration in a `config` file, the RIP parameters are whatever was saved in that file.
- If you have not saved your configuration in a `config` file, the RIP parameters are supplied by switch defaults.

This applies to all IP interfaces defined before you changed the RIP control type.

If you delete a VLAN, the parameters associated with it are removed from all associated tables. If you add a segment to a VLAN, the new segment adopts the existing RIP parameters. If you delete a segment from a VLAN, the RIP parameters remain in effect for those segments still assigned to the VLAN.

You can use the **rip-bridging enable** command to set one or more requested RIP parameters for an IP interface address and add the information to the RIP update control table. The syntax for this command is:

```
rip-bridging|rb enable|disable
```

Here is an example of the use of the **rip-bridging enable** command:

```
56:ASN-9000:ip/rip#rb enable
Okay
57:ASN-9000:ip/rip#
```

## 15.2.2 Enabling Acceptance of Default Routes

RIP is used to convey information about routes to destinations, which may be individual hosts, networks, or a special destination used to convey a default route. To enable acceptance of a default route in updates sent to your network, issue the following command:

```
ad nenable <ifaddr>
```

**<ifaddr>** Specifies the IP interface address for which you want enable acceptance of a default route. You can specify a specific IP interface address or a comma-separated list of addresses.

To disable the acceptance of default routes, issue the following command:

```
ad ndisable <ifaddr>
```

**<ifaddr>** Specifies the IP interface address for which you want to disable acceptance of a default route. You can specify a specific IP interface address or a comma-separated list of addresses.

## 15.2.3 Enabling Authentication of RIP Updates

To enable authentication of RIP Version 2 updates sent to the network, issue the following command:

```
auth nenable <ifaddr>
```

**<ifaddr>** Specifies the IP interface address for which you want to enable authentication of RIP Version 2 updates. You can specify a specific IP interface address or a comma-separated list of addresses.

To disable authentication of RIP Version 2 updates sent to the network, issue the following command:

```
auth ndisable <ifaddr>
```

**<ifaddr>** Specifies the IP interface address for which you want to disable authentication of RIP Version 2 updates. You can specify a specific IP interface address or a comma-separated list of addresses.

## 15.2.4 Setting the Authorization String on a VLAN

To set an authorization sting or a key identifier on a specified VLAN, issue the following command:

```
auth nset [-k <keyid> | <password>] <ifaddr>
```

- k<keyid>** Specifies the value to be used as the Authentication Key that has a simple password value. If a string shorter than 16 octets is supplied, the string will be left-justified and padded to 16 octets, on the right, with nulls (0x00).
- <password>** Specifies the type of Authentication used on the interface.
- <ifaddr>** Specifies the IP interface address for which you want to set an authorization sting or a key identifier on a specified VLAN. You can specify a specific IP interface address or a comma-separated list of addresses.

To unset an authorization sting or a key identifier on a specified VLAN, issue the following command:

```
auth nunset [-k <keyid> | <password>] <ifaddr>
```

- <ifaddr>** Specifies the IP interface address for which you want to unset an authorization sting or a key identifier on a specified VLAN. You can specify a specific IP interface address or a comma-separated list of addresses.

## 15.2.5 Enabling Report of Learned Routes

To enable the reporting of learned routes in updates sent to the network, issue the following command:

```
rl nenable <ifaddr>
```

- <ifaddr>** Specifies the IP interface address on which you want to enable the advent of routes in updates sent to the network.

To disable the reporting of learned routes in updates sent to the network, issue the following command:

```
rl ndisable <ifaddr>
```

**<ifaddr>** Specifies the IP interface address on which you want to enable the advent of routes in updates sent to the network.

## 15.2.6 Setting the Receive and Transmit Type on a VLAN

To set an the receive type on a specified VLAN, issue the following command:

```
rxtype nset rip1|rip2|both <ifaddr>
```

**rip1|rip2|both** Specifies if you want to specify the receive type of RIP-1 packets, RIP-2 packets, or both RIP-1 packets, RIP-2 packets on a VLAN.

**<ifaddr>** Specifies the IP interface address for which you want to set receive type on a specified VLAN. You can specify a specific IP interface address or a comma-separated list of addresses.

To set an the transmit type on a specified VLAN, issue the following command:

To set an the receive type on a specified VLAN, issue the following command:

```
txtype nset rip1|rip1c|rip2 <ifaddr>
```

**rip1|rip1c|rip2** Specifies if you want to specify the transmit the following type of RIP packets on a VLAN:

**rip1** Specifies that RIP-1 messages are sent.

**rip1c** Specifies that RIP-2 messages are sent broadcast.

**rip2** Specifies that RIP-2 messages are sent multicast.

**<ifaddr>** Specifies the IP interface address for which you want to set transmit type on a specified VLAN. You can specify a specific IP interface address or a comma-separated list of addresses.

## 15.3 Displaying and Clearing RIP Statistics

The `rip` subsystem maintains statistics on RIP packets that it transmits and receives. Use the **stats** command to display statistics. You can display statistics accumulated since the last system reset, or since the most recent statistics clear. Here is the syntax for the **stats show** command:

```
stats [show] [-t]
```

- t Displays statistics accumulated since the last switch reset. If you do not use this argument, the statistics accumulated since the last statistics clear are displayed.

Here is an example of the information displayed by this command:

```
54:ASN-9000:ip/rip# stats

RIP Packet Statistics (count since last stats clear):

RIP processing queue:

    Free entries:  150

RIP route timeout queue:

    Free entries:  2048
```

Use the **stats clear** command to clear statistics. As soon as this command is issued, the ASN-9000 clears the counters for statistics collected since the last statistics clear. Statistics accumulated since the last reboot are not cleared.

## 15.4 Bridging RIP Updates Across VLANs

By default, the ASN-9000 RIP software does not bridge RIP updates, even within an IP VLAN. Instead, the software routes the RIP updates, incrementing the metric on each route in the update by at least 1. (Routes are incremented by more than 1 if you have manually added an additional cost to an IP route listed in the switch's IP route table.)

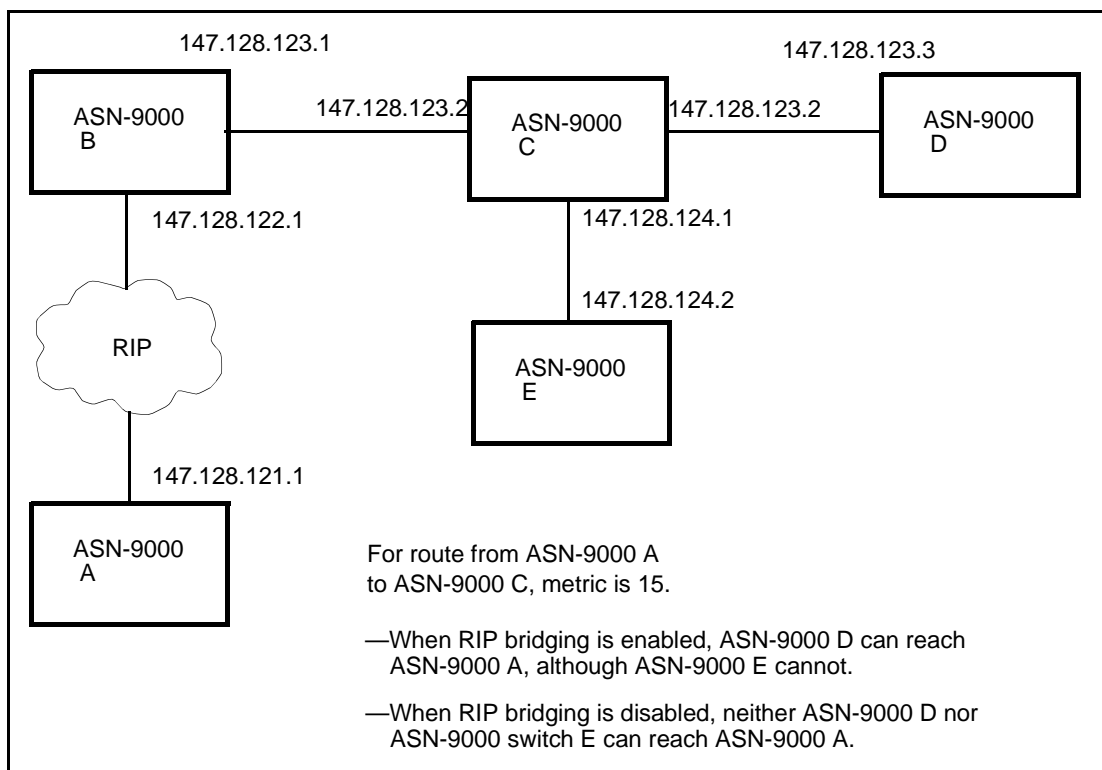
When RIP bridging is enabled, the ASN-9000 still routes RIP updates when they go from one network to another, and increments the metrics for the routes. However, the software bridges RIP updates when they go from one segment to another within a VLAN. When the software bridges a RIP update, the metrics for the routes contained in the update are not incremented.

Use the **rip-bridging enable** to configure the ASN-9000 to bridge RIP updates across IP VLANs. When the software bridges a RIP update, the metric associated with routes in the update is not incremented. Therefore, the switch does not unnecessarily add a hop to the route. Here is the syntax for this command:

**rip-bridging|rb [enable|disable]**

**enable|disable** Enables or disables the RIP bridging feature. The default is **disable**. If you omit this optional argument, the current status of the feature is displayed.

Figure 15.1 shows an example of what RIP bridging does.



**Figure 15.1 - Example of RIP Bridging**

In the above figure, two segments on ASN-9000 C have been configured with IP address 147.128.123.2. Because both segments use the same IP address, they are in a VLAN. IP traffic that normally is routed between different IP networks is bridged within the VLAN.

Suppose that the cost of the route from ASN-9000 C to ASN-9000 A is 15 hops. Because 15 hops is the maximum number of hops allowed by RIP, an additional hop would make ASN-9000 A unreachable. In the network shown above, if RIP bridging is disabled, ASN-9000 D and E cannot reach ASN-9000 A. If RIP bridging is enabled, ASN-9000 D can reach ASN-9000 A because ASN-9000 C does not increment the metric for the route to A before reporting the route to ASN-9000D.



# CHAPTER 16 IP/OSPF Subsystem Commands

This chapter lists the ASN-9000 requirements for using OSPF and describes basic features of OSPF. For complete information about OSPF, see RFC 2178. The ASN-9000 implementation of OSPF is based on this RFC.

## 16.1 Accessing the IP/OSPF Subsystem

---

To access the `ip/ospf` subsystem, issue the following command at the runtime prompt:

```
ip/ospf
```

## 16.2 Configuring an ASN-9000 Switch as an OSPF Router

---

You can configure the ASN-9000 as the following types of OSPF router:

- Internal
- Backbone<sup>1</sup>
- Area Border
- Autonomous System Border

An OSPF router can function as more than one of the router types listed above. For example, a ASN-9000 that has interfaces attached to the backbone and to other OSPF areas will function both as a Backbone router and as an Area Border router.

---

<sup>1</sup> You do not need to perform any special configuration steps to use the ASN-9000 as a Backbone router. If your configuration of the switch as an Area Border router or Autonomous System Border router places the switch on the backbone, the switch also functions as a Backbone router.

Unless you configure OSPF areas using the `area add` command, the ASN-9000 assumes that you are configuring the ASN-9000 as a Backbone router. In addition, the software automatically configures the area ID 0.0.0.0 for the backbone.

Generally, you do not need to worry about the differences among these router types. The ASN-9000 OSPF software determines how the switch is being used based upon your network configuration.

To configure the ASN-9000 for OSPF routing, you need to perform the following tasks. These tasks apply to all OSPF router types.

- Allocate memory for OSPF.
- Add IP interfaces (if interfaces are not already configured) *Chapter 13, IP Subsystem Commands*.
- Enable IP forwarding (if it is not already enabled) *Chapter 13, IP Subsystem Commands*.
- Assign the OSPF router ID.

Depending upon the type of OSPF router you plan to use the ASN-9000 as, you might need to perform some additional configuration tasks.

- If you plan to use the switch as an Interior router or an Area Border router, you need to add OSPF areas, then add OSPF interfaces to the areas.
- If your network contains areas that are not connected to the backbone and are not connected to each other, and the Area Border router for one of these areas is not a ASN-9000, you might need to create virtual links.
- If you plan to use the switch as an Autonomous System Border router, you need to enable the switch as this type of router.

Finally, after you complete the OSPF configuration steps listed above, you need to enable OSPF routing. The following sections describe how to perform these tasks.

### 16.2.1 Allocating Memory

You must allocate a portion of the ASN-9000 main memory for the `ospf` subsystem. If you do not allocate memory to the `ospf` subsystem, you cannot access it. To allocate memory for the `ospf` subsystem, issue the following command:

```
getmem
```

### 16.2.2 Assigning the OSPF Router ID

Each OSPF router within the Autonomous System must have a unique OSPF router ID. The OSPF router ID is a 32-bit address in IP format. The software does not assign an address automatically.

You can use any 32-bit address for the OSPF router ID. However, FORE Systems recommends that you use one of the IP addresses configured on the ASN-9000. By using one of the IP address on the ASN-9000, you can ensure that OSPF IDs remain unique. If you choose an IP address configured on the switch, your choice does not affect IP or OSPF. That is, the software does not establish a special relationship between the IP address you choose and the OSPF software.

By requiring that you use an IP address configured on the switch, the ASN-9000 OSPF software ensures that your OSPF router ID remains unique regardless of changes in the network. To assign the OSPF router ID, issue the following command:

```
router-id set <router-id>
```

**<router-id>** Specifies the OSPF router ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 through 9).


**NOTE**

You can define the OSPF router ID only when OSPF routing is disabled. To verify that OSPF routing is disabled, issue the **config show** command.

Here is an example of this command.

```
2:ASN-9000:ip/ospf# router-id set 1.1.1.1
```

## 16.2.3 Displaying the Router-ID

To display the router-id table, issue the following command:

```
router-id [show]
```

Here is an example of this command.

```
378:ASN-9000:ip/ospf# router-id show
OSPF Router           : memory available
OSPF Routing          : Disabled
OSPF Router ID        : NOT DEFINED!
OSPF Version number   : 2
OSPF Autonomous System Boundary Router: Enabled
Automatic Virtual Link Feature : Enabled
379:ASN-9000:ip/ospf# help router-id
```

## 16.2.4 Enabling OSPF

To enable OSPF, issue the following command:

```
ospf enable
```

To disable OSPF, issue the following command:

```
ospf disable
```

## 16.2.5 Enabling the ASN-9000 as a System Border Router

You can enable the ASN-9000 to function as an Autonomous System Border router. A switch enabled to be an Autonomous System Border router automatically exports OSPF routes to the networks outside of the OSPF Autonomous System and imports routes from the networks outside the Autonomous System. To enable or disable the ASN-9000 as an Autonomous System Border router, issue the following command:

```
asbd enable|disable
```

<b>enable disable</b>	Specifies whether you want to enable or disable the ASN-9000 to function as an Autonomous System Border router. If you specify <b>enable</b> , the ASN-9000 can exchange route information between RIP and OSPF. If you specify <b>disable</b> (disabled), the software cannot exchange route information. The default is <b>disable</b> .
-----------------------	--

To view the changes you've made, issue the following command:

```
asbd [show]
```

Here is an example of this command:

```
4:ASN-9000:ip/ospf# asbd show
```

OSPF Router:	memory available
OSPF Routing:	Enabled
OSPF Router ID:	1.1.1.1
OSPF Version number:	2
OSPF Autonomous System Boundary Router:	Enabled
Automatic Virtual Link Feature:	Enabled

## 16.2.6 Setting the Automatic Virtual-Link Feature

The automatic virtual-link feature builds virtual links between the areas that are not connected to the backbone. By building the virtual links, the ASN-9000 ensures that complete route information reaches all the OSPF routers in the Autonomous System.

### NOTE

The automatic virtual-link feature establishes virtual links only between ASN-9000's. To create a virtual link between the ASN-9000 and another type of router, use the **vlink add** command on the ASN-9000. See the documentation for your other router for information on establishing that router's end of the virtual link.

The automatic virtual-link feature is enabled by default. To disable (or re-enable) the feature, issue the following command:

```
auto-vlink enable|disable
```

**enable|disable** Specifies whether you want to enable or disable the automatic virtual-link feature.

### 16.2.6.1 Displaying the Virtual-Link Table

To display the Virtual-Link table, issue the following command:

```
auto-vlink [show]
```

Here are the results produced by this command:

```
335:ASN-9000:ip/ospf# auto-vlink
OSPF Router                : memory available
OSPF Routing                : Disabled
OSPF Router ID              : NOT DEFINED!
OSPF Version number         : 2
OSPF Autonomous System Boundary Router: Enabled
Automatic Virtual Link Feature : Enabled
336:ASN-9000:ip/ospf# help auto-vlink
```

## 16.2.7 Adding an OSPF Interface to an Area

An OSPF interface is automatically added to the ASN-9000 when you add an IP interface. The OSPF interface has the same address as the IP interface. When you enable OSPF routing, the interface is automatically added to the backbone area (0.0.0.0). To show the OSPF interface, issue the following command:

```
interface|it [show] <ip-addr>
```

**<ip-addr>** Specifies the IP address of the interface. Specify the interface in dotted decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9).



You cannot specify the TOS (Type-of-service). The ASN-9000 uses TOS 0 (zero, the IP TOS).

## 16.2.8 Using the NSET Command

In most Autonomous Systems, the ASN-9000 defaults for the OSPF interface parameters are appropriate for the Autonomous System. However, if you need to change a specific interface parameter, use the following command to do so.

```
nset <ip-addr> [ar <area-id>] [auth <key-str>]
[cost|c <cost>] [priority|p <priority>]
[xdelay|x <trans-delay>] [rint|r <rxmt-int>]
[hint|h <hello-int>] [rdint|d <rtr-dead-int>]
```

**<ip-addr>** Specifies the IP address of the interface. Specify the interface in dotted decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9).

The IP address must already be present in the ASN-9000's IP interface table before you can use it to create an OSPF interface.

**[ar <area-id>]** Specifies the OSPF area in which you are placing the OSPF interface. An OSPF interface can belong to only one area. The area must already be configured (using the **area add** command).

**[auth <key-str>]** Specifies the authentication string. For a simple password, you can specify any combination of up to eight numbers, letters, and special characters. For MD5 authentication, you can specify any combination of up to 16 numbers, letters, and special characters. The authentication string is case-sensitive.

If the area to which you are adding this interface does not require an authentication string, use empty quotation marks (“”).

**[cost|c <cost>]** Specifies the cost of using this interface. The ASN-9000 advertises the cost in Router Links Advertisements. You can specify a cost from 1 through 32. This parameter does not have a default value. The cost depends upon the wire speed of the segment on which you are adding this interface. Unless you are certain that you need to change the cost, FORE Systems recommends that you omit this argument and use the value determined by the ASN-9000.

**[priority|p <priority>]** Specifies this interface's priority during the election process for the DR (Designated Router). The interface with the highest priority number is elected as the DR. The interface with the second-highest priority number is elected as the BDR (Backup Designated Router).

You can specify a priority from 0 through 255. Priority increases from 1 (lowest) to 255 (highest). A priority of 0 (zero) makes this interface ineligible for becoming the DR. The default is 1.

If all OSPF interfaces within an Autonomous System have the same priority, the DR and BDR are elected based on the interface addresses. The interface with the highest OSPF address is elected as the DR. The interface with the second-highest OSPF address is elected as the BDR.

**NOTE**

Generally, an OSPF router has only one interface per area. If the ASN-9000 has multiple interfaces to the same area, the interface priority still applies.

**[xdelay|x <transdelay>]**

Specifies the interface transmission delay, which is the estimated number of seconds it takes to transmit a Link State Update packet over this interface. The ASN-9000 adds the transmission delay you specify to the ages of the LSAs contained in the Link State Update packets sent on this interface.

You specify a delay from **1** through **3600**. The default is **1**.

See RFC 2178 for information about choosing the transmission delay.

**[rint|r <rxmt-int>]**

Specifies the retransmission interval. The retransmission interval is the number of seconds between transmissions of LSAs to the OSPF routers adjacent to this interface. The retransmission interval also is used when transmitting Database Description and Link State Request packets.

You can specify an interval from **1** through **3600**. The default is **5**.

**[hint|h <hello-int>]**

Specifies the hello interval. The hello interval is the number of seconds between transmission of Hello packets on this interface. You can specify an interval from **1** through **65536**. The default is **10**.

**NOTE**

The hello interval (**hint**) and the router-dead interval (**rdint**) must match on neighbors. That is, the ASN-9000 values for these parameters must match the values on the ASN-9000's neighbor for these parameters. If the ASN-9000's OSPF neighbor also is a ASN-9000 system, you can ensure that the values match by accepting the defaults for these parameters. If the neighbor is not a ASN-9000, you might need to change the value on the neighbor or on the ASN-9000 so that the values on both routers match.

**[rdint|d <rtr-dead-int>]** Specifies the router-dead interval. The router-dead interval is the number of seconds the ASN-9000's OSPF neighbors should wait before declaring that the ASN-9000 (as an OSPF router) is down.

You can specify a router-dead interval from 1 through 65536. Specify an interval that is an even multiple of the Hello interval. The default is 40.

## 16.2.9 Adding an OSPF Area

When you enable OSPF routing, the ASN-9000 automatically creates an OSPF area for the network backbone. The area ID for the backbone is always 0.0.0.0.

Depending upon how you want to organize your network, you might need to add additional OSPF areas. To add an OSPF area to the ASN-9000, issue the following command:

```
area add<area-id> [<auth-type>] [stub-area-cost|sac <cost>]
```

**add** Specifies that you are adding an OSPF area to the ASN-9000.

**<area-id>** Specifies the area ID. Specify the area ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9). The area ID must be unique within the Autonomous System.

### NOTE

The area ID 0.0.0.0 is reserved for the Autonomous System's backbone and is already present in the ASN-9000.

**<auth-type>** Specifies the authentication type. You can specify one of the following:

**none | no** Specifies that the OSPF area you are adding does not use authentication.

**simple-password | sp** Specifies that a password is required for OSPF packets sent within this area.

`md5 | m` Specifies that MD5 authentication is required for OSPF packets sent within this area. See RFC 1321 for information about MD5 authentication.

The default is none (no authentication).

When you add an OSPF interface to this area (using the **interface** command), you specify the actual simple password or MD5 authentication key ID.



All OSPF routers in an area must have the same authentication type and the same authentication string. Also, all OSPF routers on a particular network should use the same authentication string.

**[stub-area-cost|sac <cost>]**

Specifies that the area is a stub area. Configuring an area as a stub area reduces OSPF overhead in the network by reducing the amount of OSPF route information flooded to the OSPF routers in the stub area.

The OSPF software does not flood external routing information (information about other Autonomous Systems) into the stub area. Internal routers in the stub area reach Autonomous Systems by using the default route to the stub area's Area Border router.

The OSPF software advertises the default route automatically. Note that a stub area's default route is unrelated to the default routes you can define in the `ip` subsystem. OSPF uses the default routes it defines in preference to manually configured default routes.

The cost is the metric for the default route out of the stub area. The stub area's Area Border router advertises the cost as part of the default route. You can specify a value from 1 through 65535. The default is 1.

### 16.2.9.1 Deleting an OSPF Area

To delete an OSPF area, issue the following command:

```
area delete|del <area-id>|all
```

**delete|del** Specifies that you are deleting an OSPF area from the ASN-9000.

**<area-id>|all** Specifies the area you want to delete. To delete all the OSPF areas defined on this ASN-9000, specify **all**.

#### NOTE

You must disable OSPF routing before deleting an area.

You cannot delete the backbone area (0.0.0.0).

### 16.2.10 Displaying an OSPF Area

Use the **area|ar [show|sh] [<area-id>]** command to display information about the OSPF areas configured on the ASN-9000.

Here are some examples of the information displayed by this command. In the following example, information is displayed for all the OSPF areas configured on the ASN-9000.

```
12:ASN-9000:ip/ospf# area show
```

Area Id	Auth Type	ImportAS ExtLSAs	Numberof SpfRuns	#Area Bdr	#AS Bdr	Number of Area LSAs	Stub Area Cost
0.0.0.0	no	Enabled	12	4	4	13	-----
1.1.1.1	no	Enabled	12	2	2	15	-----
1.2.3.4	md5	Enabled	12	0	0	0	-----
2.3.4.5	sp	Enabled	12	0	0	0	-----
3.3.3.3	no	Enabled	12	1	1	16	-----
33.0.33.0	no	Enabled	12	0	0	0	-----
33.33.33.33	no	Enabled	12	0	0	0	-----

In the following example, information is displayed for a specific OSPF area.

```
12:ASN-9000:ip/ospf# area show 1.2.3.4
```

Area Id	Auth Type	ImportAS ExtLSAs	Numberof SpfRuns	#Area Bdr	#AS Bdr	Number of Area LSAs	Stub Area Cost
1.2.3.4	md5	Enabled	12	0	0	0	----

The fields in this display show the following information:

<b>Area ID</b>	Displays the OSPF area ID you assigned using the <b>area add</b> command. The area ID is a 32-bit integer expressed in dotted decimal notation. The area ID 0.0.0.0 is the backbone area ID and is added automatically by the ASN-9000.
<b>Auth Type</b>	Displays the authentication type you assigned for this area using the <b>area add</b> command. The authentication type can have one of the following values:  no      No authentication is required for this area.  sp      A simple password is required for this area.  md5     MD5 authentication is required in this area. See RFC 1321 for information about MD5.
<b>Import AS Ext LSAs</b>	Specifies whether this area is configured to import external LSAs from other Autonomous Systems. The value can be Enabled or Disabled. To change the state of this parameter, use the <b>asbd enable disable</b> command.
<b>Number of SPF Runs</b>	Indicates the number of times the ASN-9000 has calculated this area's intra-area route table. This number is reset to zero if you disable OSPF routing, reboot the software, or power down the ASN-9000.
<b># Area Bdr</b>	Indicates the number of Area Border routers that can be reached from this area.
<b># AS Bdr</b>	Indicates the number of Autonomous System Border routers that can be reached from this area.
<b>Number of Area LSAs</b>	Indicates the number of LSAs in this area's LSA database. This number does not include external LSAs.
<b>Stub Area Cost</b>	If this area is a stub area, the metric for the stub area is indicated in this field. If this area is not a stub area, this field contains dashes (-----). You assign a stub area's metric when you add the area using the <b>area add</b> command.

## 16.2.11 Adding Network Ranges

You do not need to add network ranges to OSPF areas. The ASN-9000 automatically advertises all the networks on all the OSPF interfaces on the switch to other OSPF routers. You can add network ranges to reduce OSPF overhead or to hide certain networks from other OSPF routers.

When you add a network range to an area, link-state information for the networks within the range is summarized in the LSAs sent by the switch to its OSPF neighbors. Therefore, if you have many networks within an area, adding the networks as a network range can help reduce OSPF overhead.

In addition, you can use the **noadv** argument with the **net-range** command to prevent the switch from advertising routes to the networks within a network range. When the switch sends LSAs to its neighbors, LSAs for the networks in the hidden network range are not sent to the switch's neighbors. Therefore, other routers in the Autonomous System do not learn about the hidden networks.

### NOTE

None of the networks within the network range you add to an area can be in other areas.

To add a network range to an OSPF area, issue the following command:

```
net-range add <area-id> <net> <mask> [noadv|na]
```

- |                        |   |
|------------------------|---|
| <b>&lt;area-id&gt;</b> | Specifies the OSPF area. The area must already be added to the switch. To add an area, use the <b>area add</b> command.   |
| <b>&lt;net&gt;</b>     | Specifies an IP network address in dotted decimal notation (xxx.xxx.xxx.xxx, where each "x" is an integer from 0 – 9). The address you specify is ANDed with the subnet mask you specify for the <b>&lt;mask&gt;</b> argument.  |
| <b>&lt;mask&gt;</b>    | Specifies the IP mask associated with the IP network address you specified for the <b>&lt;net&gt;</b> argument. The mask indicates the portion of the IP network address that is to be regarded as the network portion of the address. Specify the mask in dotted decimal notation (ex: 255.255.255.0). |

<b>noadv na</b>	Prohibits the OSPF software from advertising this network range in the LSAs transmitted by the switch to its OSPF neighbors. If you use this argument, other OSPF routers do not learn about the presence of the network range.
-----------------	---

In the following example, the network range specified by IP address 200.200.200.0 and subnet mask 255.255.255.0 is added to area 1.1.1.1. When area 1.1.1.1 sends LSAs to other areas, the LSAs will contain summary information for the networks within the network range, instead of detailed link-state information for each network within the network range.

```
9:ASN-9000:ip/ospf# net-range add 1.1.1.1 200.200.200.0 255.255.255.0
OSPF: Net "200.200.200.0" with Mask "255.255.255.0" added to area "1.1.1.1"
```

If the **noadv** argument had been specified with the command, the area would not report the networks within the specified network range.

## 16.2.12 Deleting Network Ranges

To delete a network range, issue the following command:

```
net-range delete|del <area-id> <net> <mask>
```

<b>&lt;area-id&gt;</b>	Specifies the OSPF area.
<b>&lt;net&gt;</b>	Specifies the IP network address.
<b>&lt;mask&gt;</b>	Specifies the subnet mask associated with the IP address.

Here is an example of this command.

```
10:ASN-9000:ospf# net-range del 1.1.1.1 200.200.200.0 255.255.255.0
OSPF: Net "200.200.200.0" with Mask "255.255.255.0"
deleted from area "1.1.1.1"
```

After you delete a network range, the ASN-9000 sends detailed link-state information for each network, instead of summarizing the link-state information for the entire range.

## 16.2.13 Displaying Network Ranges

Use the **net-range show** [**<area-id>**] command to display information about the network ranges assigned to the areas configured on the ASN-9000. If you omit the optional **<area-id>** argument, summary information is displayed for all the network ranges in all the areas on the switch. To display network-range information for a specific area, use the **<area-id>** argument.

Here is an example of the information displayed by the **net-range show** command. In this example, the optional **<area-id>** argument is omitted. Only one network range is listed in the display, indicating that only one OSPF network range has been configured on the switch.

```
19:ASN-9000:ospf# net-range show
Area ID      Net          Mask          Advertise
-----
1.1.1.1      200.200.200.0 255.255.255.0 Enabled
```

The fields in this display show the following information:

<b>Area ID</b>	The OSPF area that contains the network range.
<b>Net</b>	The IP address of the network or subnet portion of the network range. The network number is ANDed with the subnet mask (see the Mask field) to make the network range.
<b>Mask</b>	The subnet mask that is ANDed with the network number (see the Net field) to make the network range.
<b>Advertise</b>	Indicates whether this network range is advertised to other areas. The advertise state can be Enabled or Disabled. The advertise state is enabled by default. To prevent the switch from advertising the network range to other areas, use the <b>noadv</b> argument with the <b>net-range</b> command.

## 16.2.14 Displaying OSPF Neighbors

Use the **neighbor show** command to display information about the ASN-9000 OSPF neighbors. Here is an example of the information displayed by this command.

```
18:ASN-9000:ip/ospf# neighbor show
IP Address   Router ID   Pri   State   Events   RTrQ
-----
129.213.72.2 5.5.5.5     1     full    6         0
150.1.100.3  3.3.3.3     1     full    6         0
```

The fields in this display show the following information:

<b>IP Address</b>	Neighbor's interface IP address.
<b>Router ID</b>	The ID of the OSPF router that contains the neighbor.
<b>Pri</b>	The priority of the OSPF router that contains the neighboring interface. The priority is used when the ASN-9000 elects a Designated router and a Backup

Designated router. If the priority is 0 (zero), the OSPF router is ineligible to become the Designated router or Backup Designated router.

**State** The state of the relationship with the neighboring interface's router. The state can be one of the following:

down	The switch has not received recent information from the neighbor.
attempt	The switch has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the <b>hint</b> argument of the <b>nset</b> command.
init	The switch recently received a Hello packet from the neighbor.
two Way	Communication between the switch and the neighbor now is bi-directional.
ex start	The switch and its neighbor are beginning to exchange their link-state databases.
exchange	The switch is sending its link-state database to the neighbor.
loading	The switch is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the switch sent to that neighbor. The switch updates its link-state database with the new LSAs received from the neighbor.
full	The switch and the neighbor have finished exchanging their link-state databases.

For more information about these states, see Section 10.1 in RFC 2178.

<b>Events</b>	The number of times the state of the neighbor relationship (see the State field) has changed. See Section 10.1 in RFC 2178.
<b>RTrQ</b>	The current length of the retransmission queue.

## 16.2.15 Displaying OSPF Link-State Advertisements

Use the following command to display information about a link-state database:

```
lsdb show [<lsdbid> <rid> <type> <aid>]
```

<b>&lt;lsdbid&gt;</b>	Specifies the ID of a specific LSA (link-state advertisement).
<b>&lt;rid&gt;</b>	Specifies the OSPF router ID of the router from which the link-state database was received.
<b>&lt;type&gt;</b>	Specifies the LSA type, which can be one of the following types: r: Router LSA n: Network LSA s: Summary LSA a: Autonomous System Summary LSA e: External LSA
<b>&lt;aid&gt;</b>	Specifies the ID of the area to which the LSA applies.

If you omit the optional arguments, summary information is displayed for all the LSAs present in the ASN-9000 LSA database. To display detailed information about a specific LSA, use the optional arguments.

Here are some examples of the information displayed by this command. In the first example, summary information for all LSAs in the switch's LSA database is displayed.

```
16:ASN-9000:ip/ospf# lsdb show
Area Id      Lsdb Type      Link State ID    Router ID      Sequence
-----
0.0.0.0      routerLink      1.1.1.1          1.1.1.1        -2147483552
0.0.0.0      routerLink      2.2.2.2          2.2.2.2        -2147483303
0.0.0.0      routerLink      3.3.3.3          3.3.3.3        -2147483615
0.0.0.0      routerLink      5.5.5.5          5.5.5.5        -2147483576
0.0.0.0      networkLink     80.100.1.3       3.3.3.3        -2147483635
0.0.0.0      networkLink     129.213.72.2     5.5.5.5        -2147483635
0.0.0.0      summaryLink     87.0.0.0         2.2.2.2        -2147483348
0.0.0.0      summaryLink     150.1.100.0      1.1.1.1        -2147483578
0.0.0.0      summaryLink     150.1.100.0      3.3.3.3        -2147483632
0.0.0.0      summaryLink     152.16.0.0       3.3.3.3        -2147483631
0.0.0.0      summaryLink     170.170.1.0      3.3.3.3        -2147483630
0.0.0.0      asSummaryLink   1.1.1.1          3.3.3.3        -2147483635
0.0.0.0      asSummaryLink   3.3.3.3          1.1.1.1        -2147483634
1.1.1.1      routerLink      1.1.1.1          1.1.1.1        -2147483638
1.1.1.1      routerLink      3.3.3.3          3.3.3.3        -2147483635
1.1.1.1      networkLink     150.1.100.3      3.3.3.3        -2147483646
1.1.1.1      summaryLink     44.0.0.0         3.3.3.3        -2147483640
1.1.1.1      summaryLink     80.100.0.0       3.3.3.3        -2147483640
1.1.1.1      summaryLink     80.200.0.0       3.3.3.3        -2147483640
<example truncated here for brevity>
```

The fields in this display show the following information:

<b>Area ID</b>	The OSPF area from which the LSA was received.
<b>Lsdb Type</b>	The type of LSA.
<b>Link State ID</b>	The ID of the LSA, in dotted-decimal notation. The LSA ID is determined by the type of the LSA as shown in Table 16.1.

**Table 16.1 - LSA Types**

LSA Type	LSA ID
An Internal router's LSA (routerLink).	The originating router's OSPF router ID.
A network LSA (networkLink).	The IP interface address of the network's DR (Designated Router).
A summary LSA (summaryLink).	The destination network's IP address.
An Autonomous System Border router's LSA (asSummaryLink).	The OSPF router ID of the Autonomous System Boundary router described by the LSA.
An Autonomous System Border router's external LSA (asExternalLink).	The destination network's IP address.

<b>Route ID</b>	The OSPF router from which the LSA was received.
<b>Sequence</b>	The sequence number of the LSA. The sequence number is a 32-bit signed integer. A higher sequence number indicates a more recent LSA. You can use the LSA sequence numbers to detect old or duplicate LSAs.

In the following example, detailed information is displayed about a specific LSA.

```
17:ASN-9000:ip/ospf# lsdb show 1.1.1.1 1.1.1.1 r 0.0.0.0
Detailed View
Area ID                : 0.0.0.0
Link State Database Type : routerLink
Link State ID          : 1.1.1.1
Originating Router ID  : 1.1.1.1
Sequence Number        : -2147483552
Advertisement Age       : 1503
Advertisement Checksum  : ccac
The OSPF Link State Database Advertisement: (26 per line)
00 00 02 01 01 01 01 01 01 01 01 01 80 00 00 60 cc ac 00 30 03 00 00 02 81 d5
48 02 81 d5 48 01 02 00 00 0a 03 03 03 03 96 01 64 01 04 00 00 0a
```

The fields in this display show the following information:

<b>Area ID</b>	The OSPF area from which the LSA was received.
<b>Link State Database Type</b>	The type of LSA. The LSA can be one of the following types:  routerLink Internal router LSA  networkLink Network LSA\  summaryLink Summary LSA  asSummaryLinkAutonomous System Border router LSA  asExternalLink External LSA
<b>Link State ID Area ID</b>	The ID of the LSA. The LSA ID depends upon the type of the LSA. Refer to Table 16. for LSA types and IDs.
<b>Originating Router ID Area ID</b>	The OSPF router from which the LSA was received.
<b>Sequence Number Area</b>	The sequence number of the LSA. The sequence number is a 32-bit signed integer. A higher sequence number indicates a more recent LSA. You can use the LSA sequence numbers to detect old or duplicate LSAs.

<b>Advertisement Age Area ID</b>	The age, in seconds, of the LSA.
<b>Advertisement Checksum Area</b>	The checksum for the LSA.
<b>The OSPF Link State Database Advertisement Area</b>	The contents of the LSA, in hexadecimal digits.

## 16.2.16 Enabling the Return-Code Prompt

The return-code prompt is intended primarily for automated interactions with the *ForeRunner* ASN-9000 command-line interface. To enable printing of command return codes in the next UI prompt, issue the following command:

```
rcprompt enable
```

To disable the return-code prompt, issue the following command:

```
rcprompt disable
```

## 16.2.17 Adding a Virtual-Link

Depending upon how you configure your OSPF network, it is possible for some areas to be completely disconnected from one another. Areas become disconnected from one another when they are not attached to the backbone and do not share a Border router.

The ASN-9000 can automatically link disconnected areas using the automatic virtual-link feature. This feature links together ASN-9000 configured as OSPF routers when those switches are separated from one another.

If some of the OSPF routers in your Autonomous System are not ASN-9000, you can link areas that are separated by defining a virtual link between the areas. The virtual link makes the disconnected areas virtual neighbors. LSAs from an area reach that area's virtual neighbor by travelling through a transit area. The transit area is an area between the two virtual neighbors that passes traffic between the neighbors.



You must add the transit area to your OSPF network before configuring the virtual link.

To add a virtual link, use the following command:

```
virtual-link|vlink add <aid> <router-id> [auth <key-str>]
[xdelay|x <trans-dly>] [rint|r <rxmt-int>]
[hint|h <hello-int>] [rdint|d <rtr-dead>]
```

The values and defaults for these arguments are the same as the arguments and defaults for the **nset** command.

## 16.2.18 Deleting a Virtual-Link

To delete a virtual link, issue the following command:

```
virtual-link|vlink delete|del <aid> <router-id>
```

- |                          |  |
|--------------------------|--|
| <b>&lt;aid&gt;</b>       | Specifies the area ID of the transit area. Specify the area ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9).              |
| <b>&lt;router-id&gt;</b> | Specifies the OSPF Router ID of the virtual neighbor. Specify the router ID in dotted decimal notation (xxx.xxx.xxx.xxx, where each “x” is an integer from 0 – 9). |

### NOTE

You can use the **virtual-link del** command to delete a virtual link created by the software automatically using the automatic virtual-link feature. However, if the automatic virtual-link feature is enabled, the software adds the link again. To prevent the software from adding a virtual link again, disable the automatic virtual-link feature by issuing the **auto-vlink disable** command.

## 16.2.19 Displaying Virtual-Links

Use the following command to display information about a virtual link:

```
virtual-link|vlink [show] [<aid> <router-id>]
```

If you omit the optional arguments, summary information is displayed for all the virtual links that exist between this ASN-9000 and other OSPF routers. To display detailed information about a virtual link, use the optional arguments.

Here are some examples of the information displayed by this command. In the first example, summary information is displayed. The switch in this example has only one virtual link to another OSPF router.

```
20:ASN-9000:ip/ospf# virtual-link show
Area ID      Router ID    IP Address   If State     Nbr State
-----
1.1.1.1      3.3.3.3     150.1.100.3 up           full
```

The fields in this display show the following information:

<b>Area ID</b>	The OSPF area on the local side of the virtual link.				
<b>Router ID</b>	The router ID of the OSPF router on the local end of the virtual link. (The ASN-9000 OSPF router ID.)				
<b>IP Address</b>	The IP address of the router on the remote end of the Virtual Link. Routers can have many IP addresses. This IP address is the one assigned to the remote router's segment that connects the remote router to the ASN-9000.				
<b>IF State</b>	The state of the virtual interface. The state can be one of the following: <table> <tr> <td>up</td><td>The interface can be used to send and receive OSPF route information.</td></tr> <tr> <td>down</td><td>The interface is unavailable for sending or receiving OSPF traffic. The interface's link state will be reported as down in LSAs sent from this OSPF router.</td></tr> </table>	up	The interface can be used to send and receive OSPF route information.	down	The interface is unavailable for sending or receiving OSPF traffic. The interface's link state will be reported as down in LSAs sent from this OSPF router.
up	The interface can be used to send and receive OSPF route information.				
down	The interface is unavailable for sending or receiving OSPF traffic. The interface's link state will be reported as down in LSAs sent from this OSPF router.				
<b>Nbr State</b>	The state of the virtual interface. The state can be one of the following: <table> <tr> <td>down</td><td>The switch has not received recent information from the neighbor.</td></tr> </table>	down	The switch has not received recent information from the neighbor.		
down	The switch has not received recent information from the neighbor.				

attempt	The ASN-9000 has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the <b>hint</b> argument of the <b>nset</b> command.
init	The ASN-9000 recently received a Hello packet from the neighbor.
two Way	Communication between the switch and the neighbor now is bi-directional.
ex start	The ASN-9000 and its neighbor are beginning to exchange their link-state databases.
exchange	The ASN-9000 is sending its link-state database to the neighbor.
loading	The ASN-9000 is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the switch sent to that neighbor. The switch updates its link-state database with the new LSAs received from the neighbor.
full	The ASN-9000 and the neighbor have finished exchanging their link-state databases.

In the following example, detailed information is displayed for a specific virtual link.

```
21:ASN-9000:ip/ospf# virtual-link show 1.1.1.1 3.3.3.3
Area ID                : 1.1.1.1
Router ID              : 3.3.3.3
IP Address             : 150.1.100.3
Transit Delay          : 1
Retransmission Interval : 5
Hello Interval         : 10
Router Dead Interval   : 60
Authorization Key String :
Authorization Failures  : 0
Virtual Interface State : up
Virtual Interface Events : 1
Virtual Neighbor State  : full
Virtual Neighbor Events  : 5
Virtual Neighbor Retransmission Que : 0
```

The fields in this display show the following information:

<b>Area ID</b>	The OSPF area on the local side of the virtual link.				
<b>Router ID</b>	The router ID of the OSPF router on the local end of the virtual link. (The ASN-9000's OSPF router ID.)				
<b>IP Address</b>	The IP address of the router on the remote end of the Virtual Link. Routers can have many IP addresses. This IP address is the one assigned to the remote router's segment that connects the remote router to the ASN-9000.				
<b>Transit Delay</b>	The interface transmission delay for this interface.				
<b>Retransmission Interval</b>	The retransmission interval for this interface.				
<b>Hello Interval</b>	The Hello interval for this interface.				
<b>Router Dead Interval</b>	The Hello interval for this interface.				
<b>Authorization Key String</b>	The authorization string for the interface. The authorization string is specified by the <code>&lt;key-str&gt;</code> argument of the <b>interface</b> command. If this field is blank, then no authorization string is required for this interface.				
<b>Authorization Failures</b>	The number of times another OSPF router tried to use this interface but did not supply the correct authorization string.				
<b>Virtual Interface State</b>	<p>The state of the virtual interface. The state can be one of the following:</p> <table><tr><td>up</td><td>The interface can be used to send and receive OSPF route information.</td></tr><tr><td>down</td><td>The interface is unavailable for sending or receiving OSPF route information. The interface's link state will be reported as down in LSAs sent from this OSPF router.</td></tr></table>	up	The interface can be used to send and receive OSPF route information.	down	The interface is unavailable for sending or receiving OSPF route information. The interface's link state will be reported as down in LSAs sent from this OSPF router.
up	The interface can be used to send and receive OSPF route information.				
down	The interface is unavailable for sending or receiving OSPF route information. The interface's link state will be reported as down in LSAs sent from this OSPF router.				
<b>Virtual Interface Events</b>	The number of times the state (see the Virtual Interface State field) has changed since OSPF routing was enabled.				

**Virtual Neighbor State** The state of the relationship with the OSPF router on the remote end of the virtual link. The state can be one of the following:

down	The ASN-9000 has not received recent information from the neighbor.
attempt	The ASN-9000 has not received recent information from the neighbor, but the software is attempting to contact the neighbor by sending Hello packets. You can change the Hello interval using the <b>hint</b> argument of the <b>nset</b> command.
init	The ASN-9000 recently received a Hello packet from the neighbor.
two Way	Communication between the ASN-9000 and the neighbor now is bi-directional.
ex start	The ASN-9000 and its neighbor are beginning to exchange their link-state databases.
exchange	The ASN-9000 is sending its link-state database to the neighbor.
loading	The ASN-9000 is sending Link State Request packets to the neighbor, requesting the LSAs that are more recent than the information contained in the link-state database the switch sent to that neighbor. The switch updates its link-state database with the new LSAs received from the neighbor.
full	The ASN-9000 and the neighbor have finished exchanging their link-state databases.

**Virtual Neighbor Events** The number of times the relationship with the remote end of the virtual link has changed since OSPF routing was enabled. The state is displayed in the **Virtual Neighbor State** field.

## 16.2.20 Timed Commands

In some router implementations, packet processing can affect timer execution. When multiple routers are attached to a single network, all doing broadcasts, this can lead to the synchronization of routing packets (which should be avoided). If timers cannot be implemented to avoid drift, small random amounts should be added to/subtracted from the timer interval at each firing.

## 16.2.21 Statistics Command

As soon as you enable OSPF forwarding, the *ForeRunner* ASN-9000 software begins collecting OSPF statistics. Use the **stats show** command to display statistics or the **stats clear** command to clear statistics.

## 16.2.22 Displaying OSPF Statistics

To display the OSPF statistics, issue the following command:

```
stats show
```

Here is an example of the information displayed by the **stats** command.

```
22:ASN-9000:ip/ospf# stats show
External Link-State Advertisements          : 0
Checksum of the External LSA Database        : 0
New Link-State Advertisements originated     : 105
Link-State Advertisements received           : 121
Neighbor Allocation Fails                     : 0
Link-State Advertisement Allocation Fails     : 121
Link-State Database Allocation Fails          : 121
Database Request Allocation Fails             : 0
Retransmission Allocation Fails              : 0
Acknowledge Allocation Fails                  : 0
OSPF Area Border Router                      : True
Total Authorization Failures                  : 0
Memory Usage: 3872 bytes used out of 524288 available
              6 fragments allocated, 7 total
```

## 16.2.23 Clearing OSPF Statistics

To clear OSPF statistics, issue the following command:

```
stats clear
```

Here is an example of this command.

```
23:ASN-9000:ip/ospf# stats clear
OSPF: Statistics Cleared.
```

The ASN-9000 clears the counters for the statistics and begins collecting statistics again. Statistics also are cleared if OSPF routing is disabled, the software is rebooted, or the ASN-9000 is powered down.



# CHAPTER 17

## AppleTalk Subsystem Commands

The ASN-9000 contains the `atalk` (AppleTalk) subsystem. Within this subsystem is a complete set of AppleTalk Phase-2 software and ASN-9000 for use with AppleTalk networks and internets.

You can configure the ASN-9000 to be used as an AppleTalk internet router to perform AppleTalk routing on any or all of its segments. You also can use the ASN-9000 as a local router or a backbone router, or as any combination of these types of routers.

This chapter describes the `atalk` subsystem commands you can use to perform the following tasks:

- Allocate memory for the AppleTalk subsystem.
- Enable AppleTalk routing.
- Show the current AppleTalk configuration.
- Add and delete an AppleTalk interface.
- Display the AppleTalk interface table.
- Display the AppleTalk route table.
- Display and clear the AppleTalk route cache.
- Add an and delete an AppleTalk zone.
- Display a list of configured or active (static and learned) AppleTalk zones.
- Set the aging time for entries in the AppleTalk ARP table.
- Display and clear the AppleTalk ARP table.
- Display a table of “named objects.”
- Display statistics for AARP, DDP, or ECHO packets.
- Clear packet statistics.
- Test the connectivity to another router.

## 17.1 Accessing the AppleTalk Subsystem

---

To access the `atalk` subsystem, enter the following command from the runtime command prompt:

```
atalk
```

## 17.2 Getting Started

---

To set up the ASN-9000 for AppleTalk routing, perform these steps:

1. Enable the AppleTalk subsystem:
  - Allocate memory for AppleTalk routing. (See Section 17.2.1.)
  - Enable AppleTalk routing. (See Section 17.2.1.2.)
2. Assign AppleTalk zone names to ASN-9000 segments. (See Section 17.3.1.1.)
3. Assign AppleTalk network (interface) addresses to ASN-9000 segments. (See Section 17.4.1.)
4. Save your AppleTalk configuration. (See Section 17.2.1.4.)

### 17.2.1 Enabling the AppleTalk Subsystem

Before you can use the AppleTalk subsystem, you must allocate sufficient main memory (DRAM) for the ASN-9000 to run the AppleTalk routing subsystem and enable AppleTalk routing.

#### 17.2.1.1 Allocating Memory

Before you can use the AppleTalk subsystem, you must allocate a portion of main memory (DRAM) for use by the AppleTalk subsystem. Regardless of how much main memory your ASN-9000 contains, you must allocate memory specifically for use by the AppleTalk subsystem.



FORE Systems recommends that you allocate memory for the AppleTalk subsystem immediately after you boot the ASN-9000 to ensure that the memory you request is available. For more information, see the *ForeRunner ASN-9000 Hardware Reference Manual*.

You cannot unallocate memory. To free allocated memory, make sure the configuration file does not contain a **getmem** command, then reboot the software.

To allocate memory for the AppleTalk subsystem, issue the following command:

```
getmem atalk
```

### 17.2.1.2 Enabling AppleTalk Routing

After you allocate memory, you need to enable AppleTalk routing. Use the **enable atalk** command to enable AppleTalk routing:

```
enable|disable atalk
```

<b>atalk</b>	Indicates that you are enabling (or disabling) AppleTalk routing on the ASN-9000.
<b>enable disable</b>	Specifies whether you are enabling or disabling AppleTalk routing. The default is <b>disable</b> .

Here is an example of this command:

```
4:ASN-9000:atalk# enable atalk
AppleTalk Routing: Enabled
```

### 17.2.1.3 Displaying the Current Configuration

Enter the **config show** command to verify that memory is allocated for the **atalk** subsystem and that AppleTalk routing is enabled. The command also displays the aging time for AARP (AppleTalk Address Resolution Protocol) entries. (See Section 17.5.2.)

```
5:ASN-9000:atalk# show config
AppleTalk Router: memory available
AppleTalk Routing: Enabled
AARP Aging Timer: 60 minutes
```

In this example, the display produced by the **show config** command shows the following information:

- Memory has been allocated for the AppleTalk subsystem.
- AppleTalk routing is available and enabled.
- The aging time for learned AARP entries is 60 minutes.

#### 17.2.1.4 Saving Your AppleTalk Configuration

After you verify your AppleTalk configuration, you can save the configuration using the **system savecfg** <file-name> command or the **tftp svcfg** <file-name> command, where <file-name> is the configuration file name. When you save your current configuration, the modifications you make to use the AppleTalk subsystem are available next time you reboot the *ForeRunner* ASN-9000 switch. For information about this command, see the *Hardware Reference Manual*.

## 17.3 Configuring ASN-9000 Segments for AppleTalk

---

Before the ASN-9000 can route AppleTalk packets, you must assign the appropriate zone names and network addresses to one or more ASN-9000 network ranges. Use the zone commands and interface commands to configure ASN-9000 network ranges for use with your AppleTalk networks.

### 17.3.1 The Zone Commands

The ASN-9000 uses Zone Information protocol (ZIP) to maintain a zone table that contains zone names associated with the ASN-9000 segments. Use the zone commands to add, display, or delete zone names.

#### 17.3.1.1 Adding a Zone Name

Use the **zone add** command to assign a zone name to a network range. A *zone name* is an alphanumeric string up to 32 characters in length. You can assign a different zone name to each network range, assign multiple zone names to the same network range, or assign the same zone name to multiple network ranges. Zone names are not required for non-seed segments. Moreover, for non-seed segments, the assigned zone names are not used. Assigned zone names are used for seed segments.

The zone name you assign to a ASN-9000 network range is used by the segment when it attempts to come up as a seed segment. Unless a conflict occurs over the use of the segment as a seed segment, the zone name becomes active for that segment.

You can use blank spaces in zone names. You can use blank spaces at the beginning, inside, or at the end of a zone name.

To add a zone name that contains a leading or trailing blank(s), use double quotes around the entire zone name, including the blank(s). Here is the syntax for the **zone add** command:

```
zone|zt add [-d]<netrange> <zone>
```

- [-d]** Specifies the default zone for this netrange.
- <netrange>** Optionally specifies a specific range of network addresses.
- <zone>** Specifies the zone name you want to assign to the specified netrange. A zone name is a string of 32 characters that are not case sensitive. (For example, the zone names ADMINISTRATION and administration are regarded by AppleTalk as identical.)

In the example that follows, the **zone add** command is used to assign the AppleTalk zone name Accounting to netrange 113-119.

```
8:ASN-9000:atalk# zone add 113-119 Accounting
Okay
```

Here is an example of how to add a zone name that contains a leading blank. In this example, the zone name also contains an internal blank.

```
1:ASN-9000:atalk# zone add 120 "Tony"
Okay
```

When you display AppleTalk zone names on the ASN-9000, the names that contain leading or trailing blanks are displayed with quotation marks to show the locations of the blanks.

Here is an example of how zone names that contain blanks are displayed.

```
2:ASN-9000:atalk# zone show -c
AppleTalk Zones Available for Configuration
```

Net-Range	Zones
-----	-----
120-120	tony
113-119	techsupport
101-112	*is

```
3:ASN-9000:atalk# name-table
Object Name    Object Type    Zone
ASN-9000      Router        tony
```

When you display the zone name in the Chooser on a Macintosh, the blank spaces appear in the zone name but the quotation marks are not displayed. An asterisk (\*) before the zone name indicates that this is the default zone name for this net-range.

### 17.3.1.2 Displaying the Zone Information

You can display information for configured zones or for active zones. A *configured zone* is a zone created using the **zone add** command. (See Section 17.3.1.) An *active zone* is a zone name that is actively being used on the AppleTalk internet. An active zone can be either a configured zone or a *learned zone*. A *learned zone* is a zone entry learned by the ASN-9000 software from other routers.

#### 17.3.1.2.1 Configured Zones

Use the **zone show** command to display a list of configured zone names assigned to *ForeRunner* ASN-9000 segments. Here is the syntax for this command.

```
zone|zt show [-c] <seglist> <zone> <netrange>
```

- [-c]** Specifies configured interface information. Does not specify dynamically entered interface information.
- <seglist>** Optionally specifies the segments for which you want to list the configured zone names.
- <zone>** Specifies the zone name you want to assign to the specified netrange.
- <netrange>** Optionally specifies a specific range of network addresses.

In the example that follows, the **zone show** command is used to display zone names for segments 1.1 and 2.1.

```
11:ASN-9000:atalk# zone show
Net-Range      Segments      Zones
-----
120            1.1          Test_zone
110            2.1          Test_zone
```

Use the **zone show** command to display information about active zones (both configured and learned). The **zone show** command shows the network address and the name for each currently active AppleTalk zone that is known to the ASN-9000. In addition, the table indicates whether the zone that is active on a particular network is that network's default zone. Note that configured zone names that are not in use are not listed.

In the example that follows, the `<zone-name>` argument is used to display only the networks on which the zone name “FORE Systems” is active. The asterisks (\*\*) to the left of the first network address range indicate that the zone name listed under Zone is the default zone name for that network.

```
13:ASN-9000:atalk# zone show FORE Systems
Net           Zone
**  2-2       FORE Systems
    3-3       FORE Systems
```

Each network has one and only one default zone name. However, the same zone name can be used in more than one network, and can be the default zone name in more than one network.

### 17.3.1.3 Deleting a Configured Zone

Use the `zone delete` command to delete a configured zone name from one or more segments. Here is the syntax for this command:

```
zone|zt delete <net-range> <zone>
```

<b>&lt;net-range&gt;</b>	Optionally specifies a specific range of network addresses.
<b>&lt;zone&gt;</b>	Specifies the segments from which you want to delete a configured zone. You can list individual segments, specify a range of segments, or specify <b>all</b> for all segments.

When you remove a configured zone name, the name disappears from the configured zone table. (Use the `zone show` command to display this table.)



When you use the `zone delete` command to remove a configured zone name, the change is immediately apparent in the Configured-Zone table, but does not affect zone names on interfaces that are currently up. The change can affect an interface if that interface is capable of seeding, and the segment on which the interface is defined is brought down, then back up.

Here is an example of the use of the `zone delete` command. At command prompt 14, a specific zone name (FORE Systems) is deleted from the net range 120.

```
14:ASN-9000:atalk# zone delete 120-120 FORE Systems
15:ASN-9000:atalk#
```

## 17.4 Configuring AppleTalk Interfaces

---

After you assign zone names to one or more ASN-9000 network ranges, you must then assign network addresses to each of these network ranges. Each network address consists of:

- Network address range.
- Combination of *<net>.<node>*.<sup>1</sup>
- Optionally, the default zone name.

### 17.4.1 Adding a Interface (Network Address)

Use the **interface add** command to assign a network address to one or more ASN-9000 segments.

You can assign a different network address to each net-range, or you can assign the same network address to multiple net-ranges. When you assign the same network address to more than one net-range, you create a VLAN (virtual LAN), a network that spans two or more net-ranges. A VLAN lets you increase the effective bandwidth of an AppleTalk network without creating additional network numbers. Here is the syntax for the **interface add** command:

```
interface|it add [-n] <seglist>
<seglist> <net>.<node> net[range] <x>-<y>
[-h] <seglist> <net>.<node> net[range]<x>-<y>
```

- |           |  |
|-----------|--|
| [-n]      | Specifies a non-AppleTalk passive backbone.  |
| <seglist> | Specifies the segment numbers to which you want to assign an AppleTalk network address. You can list individual segments, specify a range of segments, or specify <b>all</b> for all segments. |



To configure a segment as a non-seed segment, specify a network address range of 0-0. Do not specify a network address following the address range.

---

<sup>1</sup> In some books, this combination of net address and node address is called a “port node address,” an “AppleTalk protocol address,” or a “DDP address,” depending upon the context. This manual and other ASN-9000 documentation uses the term “network address” to refer to this combination.

If you want to create multiple non-seeding segments, you must issue a separate **interface add** command for each net. If you specify multiple segments with the same command, you create a VLAN.

To configure a segment for a non-AppleTalk (backbone) net, specify **-n**, rather than an address range. Do not specify a network address. A backbone net connects routers; nodes are not directly attached to the net.

**<net>.<node>** Specifies the network address you are assigning to the specified segment. The value you specify for **<net>** must be within the range specified by **<start-net>-<end-net>**.

For **<node>** you can specify a range from 1 through 253.



Do not use this argument if you are configuring a segment as a non-seed segment or for a non-AppleTalk (backbone) net.

Node addresses 254 and 255 are reserved AppleTalk node addresses for EtherTalk; do not use these addresses. If you attempt to use these addresses, an error message is displayed.

**net[range]** Specifies the network range you are assigning to a specified segment. You can specify a range from 1 through **65023**.

**<x>-<y>** Specifies the network ranges. For example, you can specify a network range of 113-119.

**[-h]** Specifies a hard-seed backbone.

Here are some examples of the use of the **interface add** command. In the first example, the network address range 220 through 500 is assigned to segment 2.1. The network address “220.150” indicates the specific AppleTalk node to which segment 2.1 is assigned:

```
19:ASN-9000:atalk# interface add 5 0-0
Port 5 Range 0-0 Added
Configured as non-seeding port.
```

The following example shows the command used to configure segment 2.1 as a non-seed segment. (Note that no network address range or network address is specified.)

```
18:ASN-9000:atalk# it add 2.1 220.150 net 220-500
Segment 2.1 Range 220-500 DDP Addr 220.150 Added
Configured as non-seeding interface.
```

### 17.4.2 Displaying Network Address Information

You can display information about ASN-9000 segments assigned to an AppleTalk network address using the **interface show** command. Here is the syntax for this command:

```
interface|it show [-c] <seglist> <netrange> <zone> [-a][-z]
```

- |                         |  |
|-------------------------|--|
| <b>[-c]</b>             | Specifies configured interface information. Does not specify dynamically entered interface information.  |
| <b>&lt;seglist&gt;</b>  | Specifies the segment numbers for which you want to display the AppleTalk network addresses. You can list individual segments or specify a range of segments that have AppleTalk interfaces. |
| <b>&lt;netrange&gt;</b> | Specifies the network range you are assigning to a specified segment. You can specify a range from 1 through 65023.  |
| <b>&lt;zone&gt;</b>     | Specifies the zone name for which you want to display network address information.   |
| <b>-a</b>               | Lists all configured and non-configured segments.  |
| <b>-z</b>               | Lists all configured and non-configured segments.  |

Here are some examples of the use of the **interface show** command. In the first example, no arguments are used with the command. Network address information is shown for all segments that have AppleTalk interfaces. Only two AppleTalk network addresses are assigned to ASN-9000 segments. Note that more than one zone can be associated with a segment. In Figure 17.1, three zone names are listed for segment 2.2.

A	B	C	D	E	F	G	H	
20:ASN-9000:atalk# <b>interface</b>								
Seg	DDP-Addr	Range	Type	NetCfg	Garn	From	ZoneCfg	Zone
2.1	220.150	220-220	ETH	config			config	Macintosh
2.2	2.128	2-2	ETH	garnrd	2.124		garnrd	Engineering
2.3	13.30	13-13	ETH	down	down			
2.4	128.65	128-128	ETH	unconfi	unconfig			

Figure 17.1 - interface show command details.

The table displayed by the **interface show** command shows the following information:

- A The Seg column lists the segment numbers.
- B The DDP-Addr column lists the net address for each segment to which a net address has been assigned. In this example, segments 2.1 and 2.2 are assigned AppleTalk net addresses.
- C The Range column lists the net address range assigned to each AppleTalk segment.
- D The Type column indicates the media type (in this case, “ETH,” or Ethernet).
- E The NetCfg column indicates whether the segment was a seed segment (making the *ForeRunner* ASN-9000 a seed router) for the network assigned to the segment, or learned the network information from another router in the net.

The NetCfg column indicates one of four states: config, unconfig, garnrd, or down. The initial state is unconfig. If a segment is the seed segment for a network, config soon appears under the NetCfg column. If the segment is not a seed segment, it instead relies upon another router for seed information. In such a case, when the segment has learned the network address from another router, the

state of the segment changes from `unconfig` to `garnrd`. If the segment is not configured as a seed segment and there is no other router on the network, the state remains `unconfig`.

If the state remains `unconfig`, the ASN-9000 is unable to find a seed for the segment. Check the connections joining the segment to the seed router. If the connections are working properly, the problem might be in the seed router itself.

If a segment has been configured but is attached to a router that is not turned on, or if a segment is attached to a working router but the segment has been either disabled or has not been added to a zone, the segment is listed as `down`.

If the state is `-cfg`, the segment is part of an AppleTalk VLAN and has gone down. The other segments in the VLAN might still be up.

- F** The `Garn From` column indicates the seed router from which the ASN-9000 got its configuration. If the ASN-9000 is the seed router, the `Garn From` field is blank.
- G** The `ZoneCfg` column indicates whether the interface is a seed router for the zone associated with the segment. Possible states are `config`, `unconfig`, `garnrd`, or `down`. See the descriptions for `NetCfg`.
- H** The `Zone` column lists the active zone(s) for the segment.

In the following example, the `-z` argument is used to limit the display to entries for the specified zone name (in this case, `FORE Systems`):

```
21:ASN-9000:atalk# it show -z FORE Systems
Seg  DDP-Addr  Range  Type  NetCfg  GarnFrom  ZoneCfg  Zone
---  -
2.2   2.128     2-2    ETH   garnr    2.12     garnrd   FORE Sys.
```

**NOTE**

If the interface table displays zeros under the DDP-Addr and Range columns, or “down” for the NetCfg and ZoneCfg columns, the segment may be down. If the segment is up, check if AppleTalk routing is enabled. See Section 17.2.1.2 for information on enabling AppleTalk routing.

### 17.4.3 Deleting a Network Address

Use the **interface del** command to remove an AppleTalk network address from a ASN-9000 segment:

```
interface|it del[ete] [-a] <seg-list>
```

**-a** Deletes the AppleTalk network address from a segment(s).

**NOTE**

Unless you use the **-a** argument, you must specify each segment to which a network is assigned in order to delete a network assigned to multiple segments.

**<seg-list>** Specifies the segments from which you want to delete the assigned AppleTalk network address. You can list individual segments, or specify a range of segments.

**NOTE**

If you delete an AppleTalk network address, or change or delete the zone name with which the deleted address was associated, we recommend that you wait a minimum of 15 minutes following the zone name change before re-adding the address. This time is needed by the devices in the AppleTalk internet to exchange update information about the network address and zone name changes.

Here is an example of the use of the **interface del** command. In this example, the interface table is displayed to show which interfaces are defined, then the unwanted interfaces are deleted.

```
22:ASN-9000:atalk# it show
```

Seg	DDPAddr	NetRang	Ty	NC	GarnFr	ZC	Zones
---	-----	-----	---	-----	-----	-----	-----
1.1							
1.2							
1.3							
1.4	220.150	220-230	ETH	config	220.15	config	Macintosh
1.5	2.128	2-2	ETH	garnrd	2.12	garnrd	FORE Sys.
1.6	220.150	220-230	ETH	config	220.23	config	Macintosh
1.7	220.150	220-230	ETH	config	220.23	config	Macintosh
1.8	220.150	220-230	ETH	config	220.23	config	Macintosh
1.9							
1.10							
1.11							
1.12							

```
23:ASN-9000:atalk# it del 1.4
Okay
```

In the example, the network address associated with segment 1.4 is deleted. Because the optional **-a** argument is not used, all the segments with which the network address is associated must be specified.

The following example uses the **interface delete** command with the **-a** argument to delete the same network address:

```
24:ASN-9000:atalk# interface del -a 1.6
Okay
```

When you use the **-a** argument, the network address is deleted from all segments to which it is assigned. In this example, network address 220.150, associated with segment 1.6, is deleted from segment 1.6 as well as segments 1.4, 1.7, and 1.8.

## 17.5 Using the AARP Table

The ASN-9000 uses AARP to create and maintain a table of translations between MAC-layer node addresses and AppleTalk node addresses. The AARP table enables ASN-9000 to look up the MAC-layer address of another device (node, router, and so on) based on the device's AppleTalk address. Entries in the AARP table facilitate transmission of packets from the ASN-9000 (acting as an AppleTalk router) to the devices for which MAC-layer addresses are listed. These entries are either static or learned:

<b>Static entry</b>	An entry that is created when you assign an AppleTalk network address to a ASN-9000 segment. Each time you assign a network address to a segment using the <b>interface add</b> command, the ASN-9000 automatically makes a corresponding entry in the AARP table. These entries cannot be deleted unless the corresponding network address is deleted.
<b>Learned entry</b>	An entry that the software automatically adds to the AARP table when it learns about a node address from another managed ASN-9000 or other AppleTalk router, or learns of the node address directly from one of its own segments. The ASN-9000 deletes learned entries when they are inactive for the <i>AARP aging time</i> .

For information on the AARP aging time, see Section 17.5.2. For each entry, the ASN-9000's AARP table lists the:

- DDP address of the node (also known as AppleTalk node address).
- Type of connection the segment has. There are four types of connections:

<b>Local</b>	Indicates a device is directly attached to the segment.
<b>Router</b>	Indicates the route was dynamically learned. Also indicates another AppleTalk router.
<b>Bcast</b>	Indicates the entry in the AARP table is broadcast to all devices in the network. A broadcast packet is denoted by a node address of <b>255</b> .
<b>blank</b>	Indicates a learned address, one that is added by the software. Blank entries also indicate that a node, not a router, is attached.

- MAC-layer address.
- Segment to which the node is attached.

## 17.5.1 Displaying AARP Entries

Use the **arp show** command to display the entries in the AARP table. Here is the syntax for this command:

```
arp|at [show] <seglist> <net.node>
```

**<net.node>** Specifies the network address for which you want to display AARP entries.

Here are some examples of the use of the **arp show** command. In the first example, the command is entered without an argument. The table displayed lists all AARP entries, both static entries and learned entries, for this ASN-9000.

```
25:ASN-9000:atalk# arp show
```

ARP TABLE:

DDP Address	Type	Mac Address	TTL	Segment(s)
-----	----	-----	---	-----
2.5	Local	00-00-ef-02-41-50	10	1.2
2.22		00-00-94-20-5f-82	20	1.2
2.255	BCast	09-00-07-ff-ff-ff	40	1.2
111.1	Local	00-00-ef-02-41-50	10	1.3,1.4
111.22		00-00-94-21-fd-1c	20	1.3
111.56		00-00-94-21-f2-43	20	1.4
111.255	BCast	09-00-07-ff-ff-ff	40	1.3,1.4
5.1	Local	00-00-ef-02-41-50	20	1.5
5.255	BCast	09-00-07-ff-ff-ff	40	1.5

You can specify a wildcard (\*) in place of **<net.node>**. In the following example, all DDP addresses with the net-address "2" are displayed.

```
27:ASN-9000:atalk# arp show 2.*
```

ARP TABLE:

DDP Address	Type	Mac Address	TTL	Segment(s)
-----	----	-----	---	-----
2.5	Local	00-00-ef-02-41-50	10	1.2
2.22		00-00-94-20-5f-82	20	1.2
2.255	BCast	09-00-07-ff-ff-ff	40	1.2



If the AARP table is blank, AppleTalk routing might not be enabled. Use the **config show** command to verify that routing is enabled.

## 17.5.2 Setting the AARP Aging Time

You can configure the ASN-9000 to maintain the AARP table by specifying the amount of time learned entries can remain inactive in the AARP table before being removed by the software. This time limit is the AARP aging interval and is independent of the aging time for routing table entries.

Use the **arp set age** command to set the number of minutes a learned AARP entry can be inactive before the ASN-9000 deletes it from the AARP table. Here is the syntax for this command:

```
arp set age | saa <time>
```

**<time>** Specifies the number of minutes that inactive entries remain in the AARP table. The minimum aging time is 3 minutes.

At command prompt in the example that follows, the **arp set age** command is used with the **<time>** argument to change the AARP aging time to 30 minutes.

```
28:ASN-9000:atalk# arp set age 30
ARP Age changed to 30 minutes.
```

## 17.5.3 Clearing the AARP Table

Use the **arp clear** command to clear all learned entries from the AARP table. Here is an example of the use of this command:

```
30:ASN-9000:atalk# arp clear
Okay
```

## 17.6 Displaying Route Information

Each ASN-9000 serving as a router in an AppleTalk internet uses RTMP (Routing Table Maintenance Protocol) to maintain a table of information about other AppleTalk routes throughout the internet. Use the **route show** command to display the AppleTalk route table. For each route, the route table lists the:

- Destination network address.
- Network address of the next hop (if the route is to another router).
- Segment number associated with the next hop.
- Cost (number of hops, or intermediate routers).
- State (good, suspect, or bad).

Periodically, each AppleTalk router (including other ASN-9000s serving as AppleTalk routers) broadcasts RTMP packets through each of its segments configured for AppleTalk to the other AppleTalk routers and nodes adjacent to it. As a result, each router in an AppleTalk network always has a current list of routes to the other networks. Here is the syntax for this command:

**route|rt [show] [-c|-r] [-t] [<seglist>] [<net>]**

- c|-r** Restricts the display to only directly connected entries (-c) or RTMP entries (-r).
- t** Displays the total number of entries in the route table.
- <seglist>** Specifies the segments for which you want to display route information.
- <net>** Specifies the AppleTalk net address for which you want to display route information.

Here is an example of the use of the **route show** command.

		A	B	C
		D	E	
31:ASN-9000:atalk# <b>route-table</b>				
Destination	Next Hop	Segments	Cost	State
2-2	----	1.5	0	good
3-3	2.61	1.5	1	suspect
220-220	----	1.4	0	good
774-774	2.61	1.5	1	bad

In this example, the routes for four destinations are shown:

- A** Under the Destination column is listed the network address range for each route in the routing table.
- B** The Next Hop column labels the network address of the router at the next hop. When a destination is local to the router, the next hop field contains dashes (----).
- C** The Segments column indicates the segment number through which the route can be reached.
- D** The number under Cost indicates how many hops (routers) a packet must pass through to reach the destination.

**E** The State column lists the state of the route.

A route can have one of three states: good, suspect, or bad. Approximately every 10 seconds, the ASN-9000 broadcasts an RTMP packet to each adjacent ASN-9000 to inform these ASN-9000s of active (good) routes. When the software does not receive this RTMP packet within 20 seconds, it changes the status for the routes from good to suspect.

After a route becomes suspect, the ASN-9000 waits an additional 20 seconds to receive the status packet. When the packet is received within 20 seconds, the status is changed from suspect back to good. If the packet is not received, the status changes from suspect to bad. When a route's status changes to bad, the ASN-9000 waits another 20 seconds for an RTMP packet. If the packet still is not received, the bad route is removed from the routing table.

Here is an example of the display produced if you use the `-c` argument, which displays entries only for directly connected networks:

```
32:ASN-9000:atalk# route show -c
```

Destination	Next Hop	Segments	Cost	State
2-2	-----	1.5	0	good
220-220	-----	1.4	0	good

Because the routes listed in this display are for directly connected destinations, no value appears under the Next Hop column for either route.

Here is an example of the display produced if you use the `-c` argument:

```
33:ASN-9000:atalk# route-table -c 1.4
```

Destination	Next Hop	Segments	Cost	State
220-220	----	1.4	0	good

The argument used to produce this display restricts the information to only those routes that are directly connected and are attached to segment number 1.4.



If the route table is blank, AppleTalk routing might not be enabled. Use the `config show` command to verify that routing is enabled.

## 17.7 Using the Route Cache

---

The AppleTalk route cache shows, for each segment, the most recently used destination networks. At any time, you can get an at-a-glance picture of AppleTalk routing activity in your network by displaying the AppleTalk route cache.

### 17.7.1 Displaying the Route Cache

Use the **cache show** command to display the AppleTalk route cache. AppleTalk is the syntax for this command.

```
cache [show] [<seglist>]
```

**<seglist>** Specifies the segments for which you want to display information in the route cache. If you do not specify a segment, information for all segments is shown.

Here is an example of the display produced by this command:

```
33:ASN-9000:atalk# cache show
Port 1.1:  empty
Port 1.2:  111.22,111.56
Port 1.3:  2.22
Port 1.4:  2.22
Port 1.5:  empty
Port 1.6:  empty
```

The contents of the route cache can change quite rapidly. As a result, successive **cache show** commands can give different results.

### 17.7.2 Flushing the Route Cache

The **cache clear** command removes all entries for all segments from the route cache. After the cache is flushed, new entries are again added, using the cache's "most-recently-used" algorithm. Thus, you can use the **cache clear** command to ensure that all entries displayed by a subsequent **cache show** command are fresh.

## 17.8 Displaying NBP Information

The ASN-9000 uses NBP (Name Binding Protocol) to associate names with AppleTalk network numbers, node addresses, socket numbers, and other services. With NBP, you can bind a meaningful name to any service in an AppleTalk internet. For example, you might use NBP to bind the name “Printer1” to a socket number to which a printer is attached.



The NBP table maintained by the ASN-9000 lists only the objects registered with the ASN-9000.

For each service registered with the ASN-9000, the NBP table lists the:

- Object name.
- Object type.
- Zone in which the object resides.

To display the NBP table, use the **name show** command. Here is an example of the information displayed by this command:

```
34:ASN-9000:atalk# name show
Object Name      ObjectType      Zone
PORT_220.150    Router         Macintosh
ASN-9000         Router         FORE Systems
```

A network administrator used AppleTalk NBP to name the two objects (services) “PORT\_220.150” and “ASN-9000.” Both objects are registered to this ASN-9000 as type “Router.” They belong to different zones, “Macintosh” and “FORE Systems,” respectively.

## 17.9 Displaying Statistics

During operation of your AppleTalk networks, the ASN-9000 collects statistics for AARP (AppleTalk Address Resolution Protocol), DDP (Datagram Delivery Protocol), and AEP (AppleTalk Echo Protocol) packets.

Use the **stats show** command to display statistics for AppleTalk ARP, DDP, or AEP packets. Here is the syntax for this command:

```
stats arp|ddp|echo [-t]
```

- arp|ddp|echo** Specifies the type of AppleTalk protocol for which you want to display statistics.
- t** Displays statistics collected since the most recent switch reset, rather than those collected since the most recent clear (using the **stats clear** command).

The types of statistics the ASN-9000 collects and displays depends upon the protocol type.

Here is an example of information displayed for the AARP protocol:

```
35:ASN-9000:atalk# stats show arp
ARP Statistics:
```

Requests received:	992
Replies received:	296
Invalid packets received:	0
Requests sent:	79
Replies sent:	0
Add arp entry failed:	0

Here is an example of the information displayed for the DDP protocol:

```
36:ASN-9000:atalk# stats show ddp
DDP Statistics
```

Out Requests:	93734
Out Shorts:	0
Out Longs:	93734
In Receives:	82180
Forward Requests:	63849
In Local Datagrams:	78658
No Proto Handler:	0
Out No Routes:	0
Too Short Errors:	0
Too Long Errors:	0
Broadcast Errors:	0
Short DDP Errors:	0
Hop Count Errors:	0
Checksum Errors:	0
Config Address Errors:	0
Local Range Conflicts	0
Config Zone Errors:	0
Memory Allocation Errors	0

Here is an example of the information displayed for the AEP (echo) protocol:

```
37:ASN-9000:atalk# stats show echo
Echo requests received: 39596
Echo replies received: 0
Echo requests sent: 0
```

**NOTE**

If a table displayed by the **stats** command contains all zeroes for the statistics amounts, AppleTalk routing might not be enabled. Use the **config show** command to verify that routing is enabled.

## 17.10 Clearing AppleTalk Statistics

---

To clear the statistics collected since the most recent clear, use the **stats clear** command:

```
stats clear arp|ddp|echo
```

**arp|ddp|echo** Specifies the type of AppleTalk protocol for which you want to clear statistics.

## 17.11 Testing a Network Address

---

You can use the **ping** command to test the accessibility of and round-trip delay to any AppleTalk node. This command sends an AEP (AppleTalk Echo Protocol) packet to the specified node. The AEP packet contains an instruction to the receiving device to forward the packet back to the sending ASN-9000, thus verifying receipt of the packet. To send an AEP packet, use the following command:

```
ping [-t <timeout>] [-size <pktsize>] <net>.<node>
```

**[-t <timeout>]** Optionally specifies the number of seconds the ASN-9000 waits to receive a reply packet from the specified node. The default is **15** seconds.

**[-size <pktsize>]** If you use the **<timeout>** argument, optionally specifies the size of the echo packet you want to send to the node. The packet size is measured in bytes. You can specify a packet size of 64-586 bytes. The default is **64** bytes.

**<net>.<node>** Specifies the network node to which you want to send the test packet.

The following example shows the results of the `ping` command when an AEP packet is successfully received by the sending ASN-9000:

```
39:ASN-9000:atalk# ping 220.150
220.150 is alive
```

If the target node to which you send an AEP packet is not found, or if the timeout expires before the return packet is received, an error message is displayed.

In such a case, check the route table for the network on which the specified target node resides. If the network is listed in the table, check the configuration for the target node to ensure it has learned the current network and zone-related information. If the route table and target node are okay, check the physical connections between the ASN-9000 and the target node.

# CHAPTER 18 IPX Subsystem Commands

This chapter describes the commands in the `ipx` subsystem and tells you how to use the commands to configure and manage the ASN-9000 as an IPX router. You can use the commands in this subsystem to perform the following tasks:

- Allocate memory for IPX routing
- Show the switch's IPX configuration
- Add, show, and delete IPX interfaces
- Enable IPX routing
- Show, add, and delete IPX routes
- Show or clear the IPX route cache
- Configure IPX RIP
- Add, show, and delete IPX servers
- Configure IPX helper addresses
- Show and clear IPX statistics
- Customize the IPX routing behavior

## 18.1 Accessing the IPX Subsystem

---

To access the `ipx` subsystem, enter the following command from the runtime command prompt:

```
ipx
```

## 18.2 Allocating Memory for IPX Routing

---

Before you can use the `ipx` subsystem, you must allocate a portion of the main memory for use by the `ipx` subsystem. Regardless of how much main memory your ASN-9000 contains, you must allocate memory specifically for use by the `ipx` subsystem.

**NOTE**

FORE Systems recommends memory for the IPX subsystem be allocated immediately after you boot the ASN-9000 to ensure that the memory you request is available. For more information, see the *ForeRunner ASN-9000 Hardware Reference Manual*.

To allocate memory for the `ipx` subsystem, issue the following command:

```
getmem
```

## 18.3 Showing the IPX Configuration

---

You can display the current IPX settings on the ASN-9000 by issuing the `config show` command. Here is an example of the information displayed by this command:

```
6:GE:ipx# config show
IPX Configuration:
```

```
IPX Router:                Memory Available
IPX Forwarding:            enabled
IPX Type20 Packet Forwarding: enabled
IPX Helper Feature:        enabled
Large RIP and SAP Packets: disabled
RIP broadcast timer interval: 60
SAP broadcast timer interval: 60
RIP aging timer interval:  180
SAP aging timer interval:  180
```

You can set any of the IP configuration items listed in this display.

<b>IPX Router</b>	Indicates whether main memory has been allocated for the IPX subsystem.
<b>IPX Forwarding</b>	Indicates whether IPX forwarding is enabled or disabled. The default setting is disabled.
<b>IPX Type20 Packet Forwarding</b>	Indicates that the ASN-9000 is configured to forward type-20 IPX packets. The default setting is enabled.
<b>IPX Helper Feature</b>	Indicates the setting of the IPX helper feature. When enabled, this feature allows the switch to forward unknown IPX broadcast packets.

<b>Large RIP and SAP Packets</b>	Indicates whether the ASN-9000 is enabled to forward large (greater than 576 bytes) IPX RIP and SAP packets. The default setting is disabled.
<b>RIP broadcast timer interval</b>	Indicates how often the ASN-9000 sends RIP broadcasts. The default is 60 seconds.
<b>SAP broadcast timer interval</b>	Indicates how often the ASN-9000 sends SAP broadcasts. The default is 60 seconds.
<b>RIP aging timer interval</b>	Indicates how many seconds a learned, unused IPX route can remain in the route table before it is removed by the software's aging mechanism. The default is 180 seconds, but if you choose a value other than the default, the RIP aging timer interval is always three times the RIP packet aging interval.
<b>SAP aging timer interval</b>	Indicates how many seconds a learned, unused IPX server can remain in the server table before it is removed by the software's aging mechanism. The default is 180 seconds, but if you choose a value other than the default, the SAP aging timer intervals always three times the SAP packet aging interval.

You can configure any of the IPX configuration items listed in this display. Sections in this chapter describe the commands you use to set these items.

## 18.4 Adding and Deleting IPX Interfaces

Use the **interface add** command to assign an IPX interface (sometimes referred to as a network number) to one or more ASN-9000 segments. When you add an interface, the software also makes an entry in the route table, to show that the network is directly connected to the specified segment. (See Section 18.7.1.) Here is the syntax for the **interface add** command.

```
interface|it add <segmentlist> <network>
      [mtu <mtu>] [met[ric] <metric>]
      [encap enet|802.3|802.2|snap]
```

**<segmentlist>** Specifies the segment number(s) to which you are assigning the IPX interface. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments.

**NOTE**

If you specify more than one segment number per interface, you are creating an IPX interface for a virtual LAN (VLAN). See *ForeRunner ASN-9000 Software Reference Manual* for information on IPX VLANs.

<b>&lt;network&gt;</b>	Specifies an IPX network number. Specify a hexadecimal number in the range from 1 through <b>fffffffe</b> .
<b>[mtu &lt;mtu&gt;]</b>	Specifies the maximum transmission unit (number of octets) for packets forwarded on this segment. Specify a number in the range from 576 through 1500. The default is 576.
<b>[met[ric &lt;metric&gt;]</b>	Specifies an additional cost (extra hops) of using the interface. You can specify a cost in the range 1–14. When the ASN-9000 reports this subnet using RIP, it adds the additional cost to the reported metric. The default metric is 0.

Here are some examples of the use of this command.

```
1:ASN-9000:ipx# interface add 1.2 1001 encap enet
Port 1.2, Network 0x1001, MTU 576, Cost 0, Frame type Ethernet II
Added
2:ASN-9000:ipx# it add 1.6 2002 encap enet metric 1
Port 1.6, Network 2002, MTU 576, Cost 1, Frame type Ethernet II
Added
```

The first command creates an IPX interface on segment 1. Because this interface is intended to be used as the primary route to the ASN-9000 from a router, no cost is specified.

The second command creates an IPX interface on segment 6. However, a cost has been added to this interface. The ASN-9000 RIP software adds this cost to the route when it reports it to the other routers attached to segment 6.

Here is an example of the **interface add** command used to add an IPX network to more than one ASN-9000 segment. This command creates an IPX VLAN.

```
23:ASN-9000:ipx# it add 1.1, 1.2 55ccdd55 576 802.2
Port 1.1, Network 55ccdd55, MTU 576, Frame type 802.2
Added
Port 1.2, Network 55ccdd55, MTU 576, Frame type 802.2
Added
```

### 18.4.1 Deleting IPX Interfaces

The `interface delete` command deletes an IPX interface. Here is the syntax for this command.

```
interface|it del[ete] <segmentlist>|all <network>|all
```

- <segmentlist>|all** Specifies the segment(s) from which you want to delete the network number. If you specify **all**, the network number is removed from all the ASN-9000 segments.
- <network>|all** Specifies the IPX network you want to delete. If you specify **all**, all IPX networks are deleted from the specified segment(s).

### 18.5 Displaying IPX Interfaces

You can view the network numbers assigned to segments by using the `interface show` command. Here is the syntax for this command.

```
interface|it [show] <segmentlist> <network>
```

- <segmentlist>** Specifies the segments for which you want to display IPX interface information. If you specify a list or range of segments, information is shown for only those segments that have IPX interfaces.
- <network>** Specifies the IPX network for which you want to display information.

The display includes the segment state—UP, if the segment is up, or DOWN, if you have disabled the segment or if the automatic segment-state detection mechanism has determined the segment to be down.

Here is an example of the information displayed by this command.

```
25:ASN-9000:ipx# interface show
Port Network Address  MTU  Encapsulation  State  Cost
-----
1.1    00001001          576    enet           UP     0
1.1    55ccdd55          576    802.2          UP     0
1.2    55ccdd55          576    802.2          UP     0
1.3    55ccdd55          576    802.2          UP     0
1.6    00002002          576    enet           UP     1
```

## 18.6 Enabling IPX Routing

---

After you define the IPX interfaces (see Section 18.4), you are ready to enable IPX forwarding. By enabling IPX forwarding, the IPX software can send and receive RIP and SAP updates, and respond to RIP and SAP requests from stations.

Use the following command to enable IPX forwarding: `[ipx] enable|disable`

<b>enable disable</b>	Specifies whether you are enabling or disabling IPX forwarding. The default state of forwarding is disabled.
-----------------------	--

### 18.6.1 Adding and Deleting IPX Routes

To assign the route to be used when forwarding to a particular network, use the **route add** command. Here is the syntax for this command.

```
route|rt add <network> <gw-net> <gw-addr> <seg> <hops> <ticks>
```

<b>&lt;network&gt;</b>	Specifies the destination IPX network number, a 32-bit value expressed as up to eight hexadecimal digits. Specify a number in the range from 1 through <b>fffffffe</b> .
------------------------	--

<b>&lt;gw-net&gt;</b>	Specifies the network number of the gateway (IPX router) through which packets for the destination network are to be routed. This network number must be one of the network numbers that is already configured on the segment specified by the <b>&lt;seg&gt;</b> argument. Specify a number in the range from 1 through <b>fffffffe</b> .
-----------------------	--

<b>&lt;gw-addr&gt;</b>	Specifies the IPX node number of the gateway (router) to which packets for the destination network should be forwarded. An IPX node number is actually a 48-bit MAC-layer address. Such an address is expressed in ASN-9000 commands as six hexadecimal bytes separated by hyphens.
------------------------	---

The gateway should be a device connected to a network that is directly attached to the ASN-9000 segment specified in the **<seg>** argument.

- <seg>** Specifies the ASN-9000 segment on which a packet should be forwarded to reach the specified gateway and, eventually, the specified network.
- <hops>** Specifies the number of hops to the destination, that is, how many gateways a packet must go through to reach the specified network.  
  
A hop-count of **1** corresponds to a direct connection. (Note, however, that you cannot add a route to a network that is directly attached.)  
  
The maximum number of hops is **15**; a hop-count of **16** is synonymous with “infinity” and means that the specified network is unreachable.
- <ticks>** Specifies the typical delay expected for a packet to reach its destination, measured in 55-mS “ticks.”  
  
In Ethernet and other networks with bandwidths greater than 1 Mb/s, each network is assumed to create a delay of one tick. If a route includes only such networks, the number of ticks should be set equal to the number of network segments in the route, which is the number of hops plus 1. However, routing paths that include slow, wide-area links (ex: 56 Kb/s leased lines) should have a larger number of ticks to account for the slow links.

Ticks are represented in IPX by 16-bit integers, so the practical maximum number of ticks is far less than the number that you can enter here. A statically-entered IPX route is always marked as “UP” when it is added. The route is automatically marked as “DOWN” when the corresponding segment is disabled, either manually in the `bridge` subsystem or automatically by the automatic segment-state detection mechanism.

When routing a packet to a remote network, the IPX routing software selects the route with the lowest number of ticks, regardless of whether it is a static route or a dynamic route. When two or more routes to a remote network have an equal number of ticks, the router chooses the route with the smallest number of hops.

An example of the **route add** command is shown below:

```
7:ASN-9000:ipx#rt add 008ffff9 96aabb69 0-0-99-88-88-8 2.32 3
Route to 008ffff9 via 96aabb69: added.
```

The result of this command is that packets directed to network 008ffff9 are forwarded on segment 2.3 to a gateway with address 0-0-99-88-88-88, and can expect to require a total of 2 hops and 3 ticks to reach a station on the destination network.

## 18.6.2 Deleting IPX Routes

You can completely eliminate a static route using the **route del** command. Here is the syntax for this command.

```
route|rt del[ete] <network> <gw-net> <gw-addr>
```

<network>	Specifies the destination IPX network number, a 32-bit value expressed as up to eight hexadecimal digits.
<gw-net>	Specifies the network number of the gateway (IPX router).
<gw-addr>	Specifies the IPX node number of the gateway (router).

## 18.6.3 Displaying IPX Routes

Use the **route show** command to display the IPX route table. Here is the syntax for this command.

```
route|rt [show] [-c|-r|-t] <seglist> <network>
```

-c -r -t	Restricts the display to one of the following: <ul style="list-style-type: none"><li>-c Only directly connected entries</li><li>-r Only remotely attached entries</li><li>-t Displays the total count of UP and DOWN routes.</li></ul>
<seglist>	Specifies the segment(s) for which you want to display route information.
<network>	Specifies the IPX network for which you want to display route information.

Here is an example of the display produced by this command:

```
60:ASN-9000:ipx# route show
```

Destnet	Gway-net	Gway-nodeaddr	Hops	Ticks	State	Age	Sgmts
-----	-----	-----	----	----	-----	---	-----
00001001	-----	-----	1	2	UP	---	1
00002002	-----	-----	1	2	UP	---	6
55ccdd55	-----	-----	1	2	UP	---	1
55ccdd55	-----	-----	1	2	UP	---	2
55ccdd55	-----	-----	1	2	UP	---	3
008ffff9	96aabb69	00-00-99-88-88-88	2	3	UP	---	8
054ffff9	f4f4f4f4	00-00-99-22-22-22	2	3	UP	---	4
064ffff9	f4f4f4f4	00-00-99-22-22-22	2	4	UP	---	4
011ffff9	96aabb69	00-00-99-11-11-11	2	3	UP	---	3
165ffff9	00fabcab	00-00-99-44-44-44	2	3	UP	---	9

Total no. of routes = 10 (10 UP, 0 DOWN)

This command displays the following information about IPX routes:

<b>Destnet</b>	The IPX network number of the destination network.
<b>Gway-net</b>	If the destination is not directly attached, this field contains the IPX network number of the gateway (IPX router) through which packets for the destination are to be routed.
<b>Gway-nodeaddr</b>	If the destination is not directly attached, this field contains the node address of the IPX gateway (router) through which packets for the destination are to be routed.
<b>Hops</b>	The number of gateways, including the ASN-9000, that a packet must go through to reach the destination. If a network is directly attached, the hop-count is 1.
<b>Ticks</b>	The number of 55-mS ticks that can be expected for a packet to reach its destination. If all of the network segments along the route have a bandwidth of 1 Mb/s or more, the number of ticks generally equals the number of hops plus 1. Otherwise, it is larger to account for the slower segments.
<b>State</b>	This is the state of the route; possible states are UP and DOWN. When a segment goes down, its state is updated in the interface table. All routes that use this

segment are marked DOWN in the route table, and all servers that are not accessible except through this segment are marked as DOWN in the server table.

When the segment comes back up, its state is again updated in the interface table. All routes that use this segment are marked as UP in the route table, and servers that are now accessible through this segment are marked as UP in the server table.

**Age** For dynamic routes, the number of seconds that have elapsed since this routing information was received. The Age field displays "---" for direct/static routes. For RIP entries, the Age field displays how long it has been since a routing update for the route has been received.

**Ports** Lists the segments on which packets for this destination should be forwarded.

The software does not contain a command to directly take a static route DOWN. To take DOWN a static route, use the **route delete** command to remove the route.

## 18.7 Displaying and Clearing the IPX Route Cache

---

The IPX route cache shows, for each segment, the most recently used destination networks. At any time, you can get an at-a-glance picture of IPX routing activity in your network by displaying the IPX route cache.

### 18.7.1 Displaying the Route Cache

Use the **cache show** command to display the IPX route cache. Here is the syntax for this command.

```
cache [show] <seglist>
```

**<seglist>** Specifies the segments for which you want to display information in the route cache. If you do not specify a segment, information for all segments is shown.

Here is an example of the output produced by this command. The cache displayed in this example is for a ASN-9000 containing 14 segments.

```
66:ASN-9000:ipx# cache show
IPX router cache:
Segment 1.1: empty
Segment 1.2: empty
Segment 1.3: 011ffff9, 96aabb69
Segment 1.4: f4f4f4f4, 054ffff9, 064ffff9
Segment 1.5: empty
Segment 1.6: empty
Segment 2.1: empty
Segment 2.2: 00000022
Segment 2.3: 00fabcab, 165ffff9
Segment 2.4: empty
Segment 2.5: empty
Segment 2.6: empty
Segment 3.1: empty
Segment 3.2: empty
```

**NOTE**

The contents of the route cache can change quite rapidly. As a result, successive **cache show** commands can give different results.

## 18.7.2 Clearing the Route Cache

The **cache clear** command removes all entries for all segments from the route cache. After the cache is flushed, new entries are again added, using the cache's "most-recently-used" algorithm. Thus, you can use the **cache clear** command to ensure that all entries displayed by a subsequent **cache show** command are fresh. In the following example, the route cache is flushed once and then quickly displayed two times.

```
67:ASN-9000:ipx# cache clear
IPX router cache flushed
68:ASN-9000:ipx# cache show
IPX router cache:
Segment 1.1: empty
Segment 1.2: empty
Segment 1.3: 96aabb69
Segment 1.4: f4f4f4f9
Segment 1.5: empty
Segment 1.6: empty
Segment 2.1: empty
Segment 2.2: empty
Segment 2.3: empty
Segment 2.4: empty
Segment 2.5: empty
Segment 2.6: empty
Segment 3.1: 00001fd1
Segment 3.2: empty

69:ASN-9000:ipx# cache show
IPX router cache:
Segment 1.1: empty
Segment 1.2: empty
Segment 1.3: 96aabb69, 011ffff
Segment 1.4: f4f4f4f4, 054ffff
Segment 1.5: empty
Segment 1.6: empty
Segment 2.1: empty
Segment 2.2: 00000022
Segment 2.3: 00fabcab
Segment 2.4: empty
Segment 2.5: empty
Segment 2.6: empty
Segment 3.1: 00001fd1
Segment 3.2: empty
```

## 18.8 Configuring IPX RIP and SAP Parameters

---

Earlier sections in this chapter describe how to add static entries to the IPX RIP and SAP tables maintained by the ASN-9000. However, the software contains additional RIP and SAP options you can configure:

- Whether updates on a per-segment or a per-VLAN basis are generated.
- Whether large (greater than 576 bytes) IPX RIP and SAP packets can be generated.
- Talk and listen (send and receive) settings for each interface or segment.

### 18.8.1 Setting the Control Type

You can set the RIP and SAP control type to change the RIP and SAP update mechanism. Using the `set ripsap-ctrl` command, you can configure the *ForeRunner* ASN-9000 software to generate and send a copy of each RIP and SAP packet on a per-VLAN basis instead of on a per-segment basis.

If your IPX configuration does not contain IPX VLANs, *ForeRunner* ASN-9000 performance will be the same whether you configure the software to generate updates on a per-segment basis or a per-VLAN basis. In this case, we recommend that you leave the configuration in its default state: generate updates on a per-segment basis.

However, if your configuration does include IPX VLANs, you can enhance performance by configuring the software to use the per-VLAN method for generating the RIP and SAP updates. When you change the control type to `vlan`, the software spends less time generating RIP and SAP updates, because it generates only a single update for each network, even if the network spans multiple segments. To change the RIP and SAP update method, issue the following command:

```
set ripsap-ctrl|rsct [normal|n vlan|v]
```

<b>normal n</b>	Specifies that RIP and SAP updates are generated on a per-segment basis. This is the default.
<b>vlan v</b>	Specifies that RIP and SAP updates are generated on a per-VLAN basis.

If no parameter is used with this command, the current control type is displayed.



This command affects only IPX RIP and SAP updates. It has no affect on IP RIP updates.

### 18.8.1.1 Displaying the RIP and SAP Control Type

To display the RIP/SAP control type, issue the following command:

```
ripsap-ctrl|rsct [show]
```

Here are the results produced by this command:

```
399:ASN-9000:ipx# ripsap-ctrl show
ripsap-ctrl-type:      normal
```

### 18.8.1.2 Adjusting the Interval and Aging Timers

You can adjust the following RIP and SAP timers in the *ForeRunner* ASN-9000 software.

The ASN-9000 IPX Software generates and transmits RIP and SAP updates at regular intervals. The RIP updates contain information about the IPX routes known to the ASN-9000. The SAP updates contain information about the UPX servers known to the ASN-9000.

The default interval for RIP and SAP updates is 60 seconds. Every 60 seconds, the ASN-9000 generates and transmits IPX RIP and SAP updates. Depending on whether you configured RIP and SAP updates to use the per-segment method or the network method, updates are generated for each segment or for each network.

Aging is a mechanism that periodically clears learned entries from the RIP and SAP tables. At an interval you specify (the aging interval) the ASN-9000 determines which of the learned entries in the table have not been recently used. For proper RIP and SAP reporting, the aging interval must be at least three times the duration of the broadcast interval. If an entry is not used during the specified interval, it is discarded.

A separate broadcast interval and aging timer are maintained for IPX RIP and for IPX SAP.

To set interval and aging timers for RIP, issue the following command:

```
timers set <transmit-intvl> [<rip-age>]
```

- <transmit-intvl>** Sets the RIP broadcast interval. Specify a value from 60 to 600 seconds. The default is 60 seconds.
- <rip-age>** Sets the RIP age timer. If specified, the RIP age timer value must be at least three times the value of the RIP broadcast interval. Specify a value between 180 and 1800 seconds. If unspecified, this argument defaults to three times the value of the RIP broadcast interval.

Here is an example of this command:

```
400:ASN-9000:ipx/rip# timers 100 300
```

To set interval and aging timers for SAP, issue the following command:

```
timers set <transmit-interval-time> [<aging-time>]
```

- <transmit-interval-time>** Sets the SAP broadcast interval. Specify a value from 60 to 600 seconds. The default is 60 seconds.
- <aging-time>** Sets the SAP age timer. If specified, the SAP age timer value must be at least three times the value of the SAP broadcast interval. Specify a value between 180 and 1800 seconds. If unspecified, this argument defaults to three times the value of the SAP broadcast interval.

Here is an example of this command:

```
400:ASN-9000:ipx/sap# timers 100 300
```

## 18.8.2 Setting the RIP Parameters

To enable IPX RIP sending (**talk**) or receiving (**listen**), use the **talk penable** and **listen penable** commands in the **ipx/rip** subsystem. Here is the syntax for these commands:

```
talk|ta penable <seglist>|all
talk|ta nenable <network>
listen|li penable <seglist>|all
listen|li nenable <network>
```

- <seglist>|all** Specifies the segments for which you are setting IPX RIP sending or receiving. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments. If you specify **all**, IPX RIP is enabled for all segments.

- <network>** Specifies the following:
- talk | ta** Enables the sending of RIP update packets to the specified network.
  - listen | li** Enables the learning of routes from RIP packets received from the specified network.

To disable IPX RIP **talk** or **listen**, use the **talk pdisable** and **listen pdisable** commands in the **ipx/rip** subsystem. Here is the syntax for these commands:

```
talk|ta pdisable <seglist>|all
talk|ta ndisable <network>
listen|li pdisable <seglist>|all
listen|li ndisable <network>
```

### 18.8.3 Setting the SAP Parameters

To enable IPX SAP sending (**talk**) or receiving (**listen**), use the **talk penable** and **listen penable** commands in the **ipx/sap** subsystem. Here is the syntax for these commands:

```
talk|ta penable <seglist>|all
talk|ta nenable <network>
listen|li penable <seglist>|all
listen|li nenable <network>
```

- <seglist>|all** Specifies the segments for which you are setting IPX SAP sending or receiving. You can specify a single segment, a comma-separated list of segments, or a hyphen-separated range of segments. If you specify **all**, IPX SAP is enabled for all segments.

- <network>** Specifies the following:
- talk | ta** Enables the sending of SAP update packets to the specified network.
  - listen | li** Enables the learning of routes from SAP packets received from the specified network.

To disable IPX SAP **talk** or **listen**, use the **talk pdisable** and **listen pdisable** commands in the **ipx/sap** subsystem. Here is the syntax for these commands:

```
talk|ta pdisable <seglist>|all
talk|ta ndisable <network>
listen|li pdisable <seglist>|all
listen|li ndisable <network>
```

## 18.8.4 Displaying the Configuration

To display the talk and listen (send and receive) settings for RIP and SAP updates, use the **config show** command in both the **ipx/rip** and **ipx/sap** subsystems within the **ipx** subsystem. Here is the syntax for this command.

```
config show [<seglist>] [<network>]
```

**<seglist>** Specifies the segments for which you want to display the IPX RIP and IPX SAP configurations. If no segment is specified, all RIP and SAP control table entries are displayed.

**<network>** Is the network address of the network for which you want to display RIP and SAP control table entries. If no network is specified, all RIP and SAP control table entries are displayed.

Here is an example of the display produced by this command if you selected **normal** by using the **set ripsap-ctrl** command.

```
91:ASN-9000:ipx/rip# config show
```

Segment	Talk	Listen
-----	----	-----
1.1	yes	yes
1.2	yes	yes
1.3	yes	yes
1.4	yes	yes
1.5	yes	yes
1.6	yes	yes

## 18.8.5 Setting the Parameters

The ASN-9000 IPX software advertises and receives IPX routing information using the IPX RIP. The ASN-9000 IPX software advertises and receives IPX server information using the SAP.



The RIP protocol used by IPX is different from RIP used in IP.

The commands for displaying the talk and listen (send and receive) settings for IPX RIP and SAP differ depending upon the update method used by the software:

- If the update method is per-segment, use the **penable** command in both the `ipx/rip` and `ipx/sap` subsystems within the `ipx` subsystem.
- If the update method is per-VLAN, use the **nenable** command in both the `ipx/rip` and `ipx/sap` subsystems within the `ipx` subsystem.

## 18.8.6 Equal RIP Route

To enable or disable accepting the first equal RIP route to the network, issue the following command:

```
one-rip-entry|onere enable|disable
```

## 18.9 Using the Server Table

---

The ASN-9000 IPX software stores information about NetWare file servers and other NetWare services in a data structure called the server table. The IPX routing software maintains a server table containing information that it uses when advertising services and responding to server information requests using SAP. The table contains two types of servers:

<b>Dynamic servers</b>	Learned by the switch through the SAP. IPX file servers, print servers, and other service providers advertise their existence using SAP. This information is learned by all IPX routers in the network. When an IPX station requires a service, it uses SAP to request server information from the nearest router.
<b>Static servers</b>	Configured by a system administrator, using the <b>server add</b> command. The IPX routing software always has SAP enabled, and services are always being discovered and advertised dynamically. Although the information learned through SAP is usually sufficient for good network behavior, there might be occasions in which you would like to make permanent entries in the server table. For example, you can make permanent entries in the server table

to ensure quick availability of service information after a network outage. Static service assignments can be used for this purpose.



Before you can add a server to the ASN-9000 IPX server table, you must add a route (to the IPX route table) to the server's net.

When responding to IPX stations' requests for the information on the "nearest" server of a given type, the ASN-9000 IPX software selects the server with the best route as determined from the route table, regardless of whether the server is static (added to the server table permanently by the **server add** command) or dynamic (learned through SAP). If there are equally good routes to two or more servers, the software chooses the server with the least number of hops in the server table.

### 18.9.1 Displaying the Server Table

To display the IPX servers known to the ASN-9000, issue the following command.

```
server [show] [-f|-a|-t]
<seglist> <network> <name> <type>
```

- [-f|-a|-t]** Specifies the type of entries to display:
- f Displays the entire server name, up to 48 characters. Otherwise, a maximum of 24 characters is displayed to keep the display to within an 80-character line.
  - t Displays only the total count of UP and DOWN server entries.
  - a Displays the network number and MAC address fo the next-hop gateway.
- <seglist>** Specifies the segment(s) for which you want to display route information.
- <network>** Specifies the IPX network number of the server.
- <name>** If you specify a server name here, only information that applies to the specified server is displayed.

**<type>** Specifies the type of service, either a mnemonic or a 16-bit number in the range 0 through **ffff**, expressed as up to four hexadecimal digits as shown in Table 18.1:

**Table 18.1 - Server-Type Mnemonics**

Mnemonics	Server-type(hex)
PRINT-QUEUE	0003
FILE-SERVER	0004
JOB-SERVER	0005
PRINT-SRVR	0007
ARCHIVE-SVR	0009
REM-BRIDGE	0024
ADVRT-PRINT	0047

Here is an example of the output produced by this command.

```
73:ASN-9000:ipx# server show
SrvrTyp  SrvrNet  SrvrNode      Sock  Hop  State Sgmt  Age  Srvr-name
-----
00ff     f4f4f4f4 00-00-99-66-66-66 f4f4  2    UP    4      4      crp-srvr
00fe     f4f4f4f4 00-00-99-66-66-66 f4f4  2    UP    4      4      big-boss
0123     f4f4f4f4 00-00-99-66-66-66 f4f4  2    UP    4      4      cat-mkr
Total no. of servers = 3 (3 UP, 0 DN)
```

This command displays the following information from the server table:

<b>Server-type</b>	Specifies the type of service, either a mnemonic or a 16-bit number in the range 0 through <b>ffff</b> , expressed as up to four hexadecimal digits.
<b>Srvr-net</b>	The IPX network number of the server.
<b>Server-node</b>	The IPX node number of the server.
<b>Sock</b>	The IPX socket number on which the server accepts requests for service.
<b>Hop</b>	The number of gateways, including the <i>ForeRunner</i> ASN-9000 switch, that a packet must go through to reach the server. If the server is on a directly-attached network, the hop-count is 1.

<b>State</b>	This is the state of the server; possible states are “UP” and “DOWN.”
<b>Segment</b>	The segment on which the entry was learned.
<b>Age</b>	For dynamic servers, the number of seconds that have elapsed since this information was received.
<b>Server-name</b>	The name of the server, up to 48 ASCII characters.

## 18.9.2 Adding a Static Server

To add a server to the server table, use the **server add** command. Here is the syntax for this command.

```
server add <s-type> <s-net> <s-addr> <s-sock> <s-hops> <s-name>
```

**<s-type>** Specifies the type of service, either a mnemonic or a 16-bit number in the range 0 through **ffffe**, expressed as up to four hexadecimal digits (refer to Table 18.1 for server types).

**<s-net>** Specifies the IPX network on which the server resides, a 32-bit number expressed as up to eight hexadecimal digits.

Note that the *ForeRunner* ASN-9000 software does not accept the **server add** command if there is no known route to the server’s network at the time the command is given. Specify a number in the range from 1 through **ffffffe**.

**<s-addr>** Specifies the IPX node number of the server. This is a 48-bit MAC-layer address, expressed as six hexadecimal bytes separated by hyphens.

**<s-sock>** Specifies the IPX socket number on which the specified server accepts requests for service.

**<s-hops>** Specifies the number of hops to the specified server, that is, how many gateways a packet must go through to reach it. The maximum number of hops is 15; a hop-count of 16 is synonymous with “infinity” and means that the specified server is unreachable.

**<s-name>** Specifies the name of the server, up to 48 ASCII characters. Server names are case sensitive.

Here is an example of the **server add** command:

```
3:ASN-9000:ipx# server add 4 fabcab 0-0-88-88-88-88 1010 2 phsrvr
Server phsrvr of type 0004 on net 00fabcab: added.
```

### 18.9.3 Deleting a Static Server

You can completely eliminate a static server assignment using the **server delete** command. Here is the syntax for this command.

```
server del[ete] <s-type> n[ame] <s-name>
```

- |                       |  |
|-----------------------|--|
| <b>&lt;s-type&gt;</b> | Specifies the type of service, either a mnemonic or a 16-bit number in the range 0 through <b>ffffe</b> , expressed as up to four hexadecimal digits (refer to Table 18.1 for server types). |
| <b>&lt;s-name&gt;</b> | Specifies the name of the server, up to 48 ASCII characters. Server names are case sensitive.  |

Here is an example of the use of this command.

```
72:ASN-9000:ipx# server del 4 eng-server
Server eng-server of type 0004: deleted from table.
```

## 18.10 Using IPX Helper

---

This section describes how to use the IPX Helper feature. IPX Helper lets the ASN-9000 forward unknown IPX broadcast packets, which normally would be dropped, onto specified networks. This feature forwards the unknown IPX broadcast packets without using the IPX SAP protocol.

When you assign an IPX helper address to a segment, and an unknown IPX broadcast packet with the specified destination socket number is received on that segment:

- The IPX broadcast packet destination network number and destination node address are replaced with the number and address specified in the **helper add** command.
- The IPX broadcast packet then is forwarded onto all other segments.

To use IPX Helper, you first must enable it by issuing the following command:

```
enable|enl disable|dis helper
```

### 18.10.1 Adding an IPX Helper Address

Use the **helper add** command to add an IPX helper address to a segment. Here is the syntax for this command:

<b>helper add</b> <seglist> <network> <node address> <socket>	
<seglist>	Specifies a segment, a comma-separated list of segments, or a hyphen-separated range of segments.
<network>	Specifies a network number or the value <b>ffffff</b> to specify all net broadcast.
<node address>	Specifies the unicast address or the broadcast address <b>ff-ff-ff-ff-ff</b> .
<socket>	Specifies a socket number in hexadecimal notation. To specify any socket number, enter the value <b>fff</b> .

Here is an example of how to add an IPX Helper address. In this example, a broadcast address is defined.

```
95:ASN-9000:ipx# helper add aabbccdd ff-ff-ff-ff-ff ffff 1
```

### 18.10.2 Displaying an IPX Helper

Use the **helper show** command to display IPX helper addresses assigned for all segments. Here is an example of the information displayed by this command.

```
220:ASN-9000:ipx# helper show
```

SEGMENT	NETWORK	NODE ADDRESS	SOCKET NUMBER
-----	-----	-----	-----
1	aabbccdd	ff-ff-ff-ff-ff	ffff

### 18.10.3 Deleting an IPX Helper Address

Use the **helper delete** command to delete an IPX helper address assigned to a segment. The syntax for this command is:

<b>helper delete</b> <seglist>	
<seglist>	Specifies a segment, a comma-separated list of segments, or a hyphen-separated range of segments.

## 18.11 Showing and Clearing Statistics

---

Use the **stats** command to display IPX or type-20 packet statistics. Here is the syntax for this command:

```
stats [show] [-t]
```

- t** Optionally displays statistics collected since the most recent switch reset, rather than those collected since the most recent clear.

Here is an example of the output produced by the **stats** command.

```
80:ASN-9000:ipx# stats
IPX statistics: count since last stats clear
Datagrams received:          2302091
Header errors received:      0
Address errors received:     0
Datagrams forwarded:         2302091
Unknown Broadcast packets forwarded: 0
Unknown protocols received:  0
Incoming datagrams discarded: 0
Datagrams delivered to higher layer: 2258
Datagrams sent:              6658
```

Here is an example of the output produced by the **stats type20** command.

```
81:ASN-9000:ipx# t20stats
Type-20 statistics: count since last stats clear
Packets received:          0
Packets forwarded:         0
Packets discarded:         0
Packets in error:          0
```

Here is an example of the use of the **-t** argument with the **stats** command. In this example, IPX statistics collected since the last switch reset are displayed.

```
83:ASN-9000:ipx# stats -t
IPX statistics: Total count since last system reset
Datagrams received:          2305309
Header errors received:      0
Address errors received:     0
Datagrams forwarded:         2305309
Unknown Broadcast packets forwarded: 0
Unknown protocols received:  0
Incoming datagrams discarded: 0
Datagrams delivered to higher layer: 2261
Datagrams sent:              6664
```

To clear statistics, use the **stats clear** command.

# 18.12Customizing the IPX Configuration

To enable or disable the forwarding of type-20 packets for the entire switch, issue the following command:

```
enable|disable type20-forwarding|t20fw
```

## 18.12.1 Type-20 Forwarding for Segments

Use the `type20-port-forwarding` command to show whether type-20 packet forwarding is enabled or disabled on specific segments. The syntax for this command is:

```
penable|pdisable type20-port-forwarding|tpfw <seglst>
```

- penable|pdisable

Specifies whether you are enabling or disabling type-20 packet forwarding. The default `type20-port-forwarding` is enabled.
- <seglst>

Specifies a segment, a comma-separated list of segments, or a hyphen-separated range of segments for which you want to enable or disable type-20 packet forwarding.

## 18.12.2 Enabling Large Packets

In software version 3.0, IPX RIP and IPX SAP packets larger than 576 bytes (the default minimum) can be generated. To change the default, use the `enable large-rip-sap-pkt` command to enable the software to generate large RIP and SAP packets. The syntax for this command is:

```
enable|disable large-rip-sap-pkt|lpkt
```

- enable|disable

Specifies that the ASN-9000 generate or not generate IPX RIP or IPX SAP packets larger than 576 bytes. The default for `large-rip-sap-pkt` is disabled.



The MTU setting for the IPX interfaces you define on the switch needs to be more than 576 bytes to generate larger RIP and SAP packets.



# CHAPTER 19

## Configuring IPX Translation Bridging

IPX translation bridging configures one or more IPX networks that span across Ethernet segments using different packet encapsulations. Without altering the configurations of individual devices, IPX translation bridging enables Ethernet devices with different encapsulation types to communicate with each other. This feature is especially useful if the IPX network consists largely of Ethernet devices using 802.3 encapsulation, the default encapsulation type in Novell IPX software versions 2.2 through 3.11. However, if the network name is not in the IBT table, IPX translation bridging does not occur, and normal bridging does. This section describes how to perform the following tasks:

- Show the switch's IPX translation-bridging configuration
- Add, show, and delete IPX translation-bridging interfaces

### NOTE

IPX translation bridging is independent of IPX routing—they are mutually exclusive. We recommend that you do not enable both IPX translation bridging and IPX routing. However, if both IPX translation bridging and IPX routing are enabled, IPX routing takes precedence over IPX translation bridging.

## 19.1 Encapsulation Types

When using IPX translation bridging, specify the Ethernet encapsulation types to be used on each IPX network. For each IPX network number, specify both the Ethernet encapsulations to be use on that network. Table 19.1 lists the combinations of encapsulation types that can be specified.

**Table 19.1 - Encapsulation types**

	ENET	802.2	802.3	SNAP
Ethernet	4	4	4	4

## 19.2 Configuration Requirements

---

Although IPX translation bridging is simple to configure, the following conditions must be met:

- The servers attached to the segments in an IPX translation bridging network must be configured to have the same network number as the “IPX translation-bridging” network number configured on the ASN-9000. If a server’s network number cannot be changed to correspond to the IPX translation-bridging network defined on the switch, change the network number to match the server.
- Servers and clients must be configured to have the same encapsulation type as the type specified for the appropriate medium in the IPX translation-bridging network. For example, a client attached to an Ethernet segment must be configured to use the same Ethernet encapsulation type as the one defined for the corresponding IPX translation-bridging network. However, if encapsulation types on the server or client cannot be changed, the encapsulation types of the client or server can be configured on the switch.

### 19.2.1 Enabling IPX Translation Bridging

Before using the IPX translation-bridging feature, IPX translation bridging must be enabled. To enable IPX translation bridging, issue the following command:

```
ipx-br-translation|ibt enable|disable
```

<b>enable disable</b>	Specifies whether to enable or disable IPX translation bridging.
-----------------------	--

Here is an example of the use of this command:

```
1:ASN-9000:bridge# ipx-br-translation enable
IPX translation bridging is now enabled
```

### 19.2.2 Adding IPX Translation-Bridging Interfaces

To create an IPX translation-bridging network, use the following command:

```
ipx-br-translation|ibt add <network> <ether-encap>
```

<b>&lt;network&gt;</b>	Specifies the IPX network number to apply the encapsulation settings.
<b>&lt;ether-encap&gt;</b>	Specifies the encapsulation type to be used for Ethernet packets. Specify one of the following:

enet Ethernet Type II encapsulation.

802.3 Raw 802.3 encapsulation.

802.2 802.3 with an LLC header.

snap 802.3 with LLC and SNAP headers.


**NOTE**

The default Ethernet encapsulation type used in Novell IPX versions 2.2 through 3.11 is 802.3. The default for versions 3.12 through 4.x is 802.2.

Here are some examples of how to use this command. In these examples, definitions are created for IPX translation-bridging networks 100, 200, and 300:

```
2:ASN-9000:bridge# ipx-br-translation add 100 802.2 snap
IPX network 100 added to the translation table
3:ASN-9000:bridge# ipx-br-translation add 200 802.2 802.2
IPX network 200 added to the translation table
4:ASN-9000:bridge# ipx-br-translation add 300 802.3 snap
IPX network 300 added to the translation table
```

### 19.2.3 Displaying IPX Translation-Bridging Interfaces

At any time, you can display the definitions for the IPX translation-bridging networks defined on the ASN-9000. To display the definitions, use the following command:

```
ipx-br-translation|ibt show [<network>] | [-t]
```

**<network>** Specifies an IPX translation-bridging network number.

**-t** Displays only the total number of entries in the IPX translation-bridge table.

Here are some examples of displays produced by this command. In the first example, no specific network number is given, so all individual entries are displayed, as well as the total number of entries. In the second example (prompt 7), the **-t** argument is used to display the total number of IPX translation-bridging entries.

## Configuring IPX Translation Bridging

```
6:ASN-9000:bridge# ipx-br-translation show
IPX Translation Bridging: Enabled
IPX Network      Ethernet Encap
-----
      100        802.2
      200        802.2
      300        Ethernet II

Total entries: 3
7:ASN-9000:bridge# ipx-br-translation show -t
IPX Translation Bridging: Enabled
IPX Network      Ethernet Encap
-----
Total entries: 3
```

### 19.2.4 Deleting IPX Translation-Bridging Interfaces

To delete the encapsulation settings assigned to a network number, use the following command:

```
ipx-br-translation|ibt del <network>|all
```

**<network>|all** Specifies an IPX translation-bridging network number. If **all** is specified, all IPX translation-bridging networks are deleted.

Here is an example of the use of this command:

```
8:ASN-9000:bridge# ipx-br-translation del all
All IPX networks deleted from the IPX translation table
```

The ASN-9000 contains a complete set of DECnet Phase IV routing software for use in DECnet networks. The routing engine works side-by-side with the Ethernet bridging software. With appropriate configuration, the ASN-9000 can be set up to perform DECnet routing on any segments.

This chapter assumes that a familiarity with the basic requirements of DECnet networks and the DECnet protocol. For further information on this subject, refer to a DECnet guide, such as the DECnet Phase IV General Description, Order No. AA-N149A-TC, (Digital Equipment Corporation, 1982).

This chapter describes the commands and facilities of the DECnet subsystem. To set up the ASN-9000 for DECnet routing. The following steps must be performed:

1. Allocate memory for DECnet routing.
2. Assign the DECnet node ID using the **set node-id** command.
3. If the ASN-9000 is to be a Level-2 router, select it with the **set node-type** command.
4. Turn on DECnet routing with the **enable dec** command.
5. Enable DECnet routing on the desired segments with the **penable dec** command.

A large number of nodes may necessitate increasing the maximum limits for these parameters with the **set max-area-num** and **set max-node-num** commands.

After setting up DECnet routing, check connectivity to hosts and other routers using the ASN-9000 **show** and **stats** commands.

After DECnet is configured, it is recommended that the configuration be saved using the **system** or **tftp savecfg** command. Refer to *Chapter 6, System Subsystem Commands* or *Chapter 10, TFTP Subsystem Commands* for details on these commands.

## 20.1 Accessing the DECnet Subsystem

---

Access the DECnet subsystem by entering **dec** at a runtime prompt. Most of the commands in this chapter assume that the command prompt is **dec**. Those commands, such as **getmem**, which are not located in the **dec** subsystem are identified by listing the subsystem name with the command (ex: **atalk getmem**.)

## 20.1.1 Allocating Memory

Before using the `dec` subsystem, memory must be allocated by issuing the `getmem` command, as shown in the following example:

```
1:ASN-9000:dec# getmem
Memory allocated for DEC routing.
2:ASN-9000:dec#
```

If memory has been allocated for DECnet routing at the time the configuration is saved, the corresponding `getmem` command is placed in the configuration file ahead of the other DECnet configuration commands. Thus, the `getmem` command only needs to be entered when first configuring the ASN-9000 for DECnet routing.



FORE Systems recommends that memory for the DECnet subsystem be allocated immediately after booting the ASN-9000 to ensure that the memory is available.

Verify that memory has been allocated using the `rs` command. The command can be executed if memory has not been allocated.

```
3:ASN-9000:dec# rs
DECnet routing status:
```

Node/Segment	Management State	Routing State
-----	-----	-----
DECnet-Forwarding	Disabled	Down
Segment 1.2	Disabled	Down
Segment 2.1	Disabled	Down
<i>&lt;additional rows omitted for brevity&gt;</i>		

```
4:ASN-9000:dec#
```

## 20.1.2 Node Configuration

When placed in a DECnet internetwork, the ASN-9000 acts as a standard router, capable of connecting many different DECnet networks together. It determines the identity and location of neighbors through standard DECnet Phase IV protocols, and finds the closest path to each. It then uses this information to route packets that arrive at the input segments.

The DECnet Phase IV routing protocol calls for each node to have an “area number” (between 1 and 63), as well as a node ID (from 1 to 1023). For Level-1 networks (consisting only of Level-1 endnodes and routers), the area numbers are identical and unused. In Level-2 networks, an “area” is defined as a collection of several nodes with identical area numbers. These areas are connected by Level-2 routers. If the ASN-9000 is configured as a Level-2 router (with the `set node-type area-router` command), an extended set of routing protocols is used that can

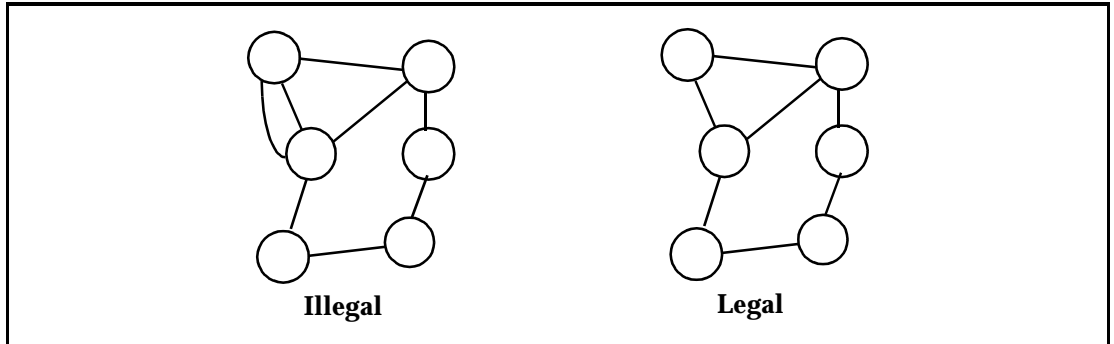
connect nodes from different areas. Normal nodes can only route packets directly to other nodes within their area (those with matching area numbers). If they are called upon to send a packet to a node in another area, they send it to the “nearest” Level-2 router. This Level-2 router keeps track of routes to all other Level-2 routers, as well as routes to normal nodes within its area. This two-level hierarchy allows for a larger network with manageable routing tables.

When the ASN-9000 is configured as a Level-2 router, it locates all nodes in its area *and* all other Level-2 routers. Note that this places some restrictions on the topology of the network, as described below. The ASN-9000 announces itself as a Level-2 router to the normal nodes in its area so that all inter-area packets are sent through it (thus a pair of Level-2 routers are needed for inter-area packets: one in each area). If the ASN-9000 is placed into a network that uses Level-2 routing, and the ASN-9000 is to serve as a Level-2 router, be sure to turn this option on (with **set node-type area-router**). If the ASN-9000 is not to be a Level-2 router, or if the network uses only Level-1 routing, turn this option off (with **set node-type router**). The default is to use only Level-1 routing.

The use of areas in DECnet Level-2 routing places some restrictions upon the topology of these networks:

- Each node must be able to get to each other node in its area without the use of Level-2 routers and without leaving its area. Consequently, all the nodes in a given area must form a contiguous group. If all nodes from other areas are removed, leaving only the nodes from this area, there can be no isolated nodes remaining. This restriction also applies to Level-2 router nodes.
- The set of all Level-2 router nodes must form a contiguous group so that any packet going from one Level-2 router to another can travel only through other Level-2 router nodes.
- There can not be multiple links between adjacent routers. If two routers are directly connected by more than one segment, the DECnet protocol must be enabled and running on only one of those links. Failure to ensure DECnet is running on only one link results in changes to the routing table every time the doubly-connected nodes discover each other. Such a double connection causes the routing table to be continually flushed, resulting in poor performance and unreachable nodes.

This situation is represented graphically in Figure 20.1:



**Figure 20.1 - Illegal Double Links**

There is also a topological consideration that improves the efficiency of DECnet Level-2 networks. When a heavily populated broadcast medium is used, all the nodes on the same segment should be assigned the same area number. The reason is that two nodes with different area numbers must use Level-2 routing to communicate. Therefore this cable segment must have a pair of Level-2 routers on it (one for each area), and the communication path requires three hops, even though the nodes are on the same segment and could communicate directly by other protocols. To avoid these extra hops, all nodes that can communicate directly with each other should be placed in the same area by giving them identical area numbers.

### 20.1.3 DECnet Network Topology Restrictions

- All nodes in a given area must be connected.
- All Level-2 nodes must be connected.
- No redundant paths are allowed between adjacent routers.

Note that these restrictions do not prevent the same network segment from serving both Level-1 nodes and Level-2 nodes. Thus the same segment can serve to connect Level-1 routers, Level-1 endnodes, and Level-2 routers. The requirement is that all of an area be contiguous; nodes from different areas can be on the same segment as long as data moving within one area does not have to pass through the other area's nodes in order to reach its destination.

### 20.1.4 Configuring the ASN-9000 as a DECnet Node

First, set the maximum node number used in this area. To do this, use the **max-node-num** parameter:

```
set max-node-num|mn <value>
```

This determines the number of nodes that can exist within the ASN-9000 area. The routing software ignores any packets from nodes outside this range. The default is 255, so you must increase it if you have nodes with larger numbers. The DECnet protocol requires node numbers to be in the range 1 to 1023, so you cannot raise the **max-node-num** parameter above 1023.

```
171:ASN-9000:dec# set mnn 1023
Okay
```

Assign the ASN-9000 node ID using the **node-id** parameter command:

```
set node-id|nid <area>.<node>
```

This command instructs the ASN-9000 to use the specified address for all DECnet communications. The **<area>** parameter must match the area in which the ASN-9000 has been placed; recall that the DECnet definition of “area” is the set of nodes that have the same area numbers. The **<node>** parameter can be any value that is unique among all nodes in the specified area.

```
172:ASN-9000:dec# set nid 5.1023
Okay
```

Select the type of routing that needs to be done by this router. Use the **node-type** parameter command:

```
set node-type|ntp router|rt | area-router|a
```

This command determines what kind of routing the ASN-9000 performs. If you choose “**router**,” the ASN-9000 performs only Level-1 routing. A Level-1 router keeps track of nodes within its own area only, and does not try to determine routes to other areas. If it receives data for another area, it sends it to the nearest Level-2 router. The ASN-9000 acts as a Level-1 router by default.

If you choose “**area-router**,” the ASN-9000 also performs Level-2 routing. Level-2 is a superset of Level-1: the node routes data to nodes within its area, as well as find routes to other areas. All Level-2 routers find all other Level-2 routers (including those in other areas), and inter-area traffic is sent to a distant Level-2 router for local distribution.

By default, the router performs only Level-1 routing. No changes need to be made to this parameter if the ASN-9000 is going to be used as a Level-1 router. For Level-2 routing, type: **set node-type area-router**.

```
173:ASN-9000:dec# set nt area-router
Okay
```

Activate DECnet routing with the **enable dec** command:

```
enable dec
```

This is the primary command which turns on all of the DECnet routing software. However, to have a useful configuration, you must still specify two or more segments that use DECnet. This is accomplished with the **penable dec <seglist>** command, described in.

```
191:ASN-9000:dec# penable dec
Okay
```

Verify the node configuration with the **show node-type** command.

```
195:ASN-9000:dec# show node-type
DECnet node configuration
-----
DEC-forwarding:      Enabled
Max-Area-Num:        63
Max-Node-Num:        1023
Max-Adj-Endnodes:    1023
Max-Adj-Routers:     128
Max-Cost-To-Area:    100
Max-Hops-To-Area:    16
Max-Cost-To-Node:    125
Max-Hops-To-Node:    30
Max-Visits:          60
Node-Type:           Area Rtr
Node-ID:             5.1023
Routing-State:       Up
Update-Time:         60 seconds
```

Now configure one or more segments to use DECnet forwarding.

### 20.1.4.1 Additional Node Commands

The following additional commands are available to set or show node parameters.

```
max-adj-endnodes|mae set <value>
max-adj-endnodes|mae [show]
```

**Sets, or displays,** the number of endnode adjacencies supported by this router. The range for *<value>* is 1 - 1023.

```
max-adj-routers|mar set <value>
max-adj-routers|mar [show]
```

**Sets, or displays,** the number of broadcast router adjacencies supported by this router. The range for *<value>* is 1 - 560.

```
max-area-num|man set <value>
max-area-num|man [show]
```

**Sets, or displays,** the maximum area number allowed in the entire network. The range for *<value>* is 1 - 63. *<value>* must be greater than or equal to the maximum area in use.

```
max-cost-to-area|mca set <value>
max-cost-to-area|mca [show]
```

**Sets, or displays,** the maximum cost possible in a path to a reachable area. The range for *<value>* is 1 - 1022. *<value>* must be greater than or equal to actual max hops to an area \* 25.

```
max-cost-to-node|mcn set <value>
max-cost-to-node|mcn [show]
```

**Sets, or displays,** the maximum cost possible in a path to a reachable node. The range for *<value>* is 1 - 1022. *<value>* must be greater than or equal to actual max hops in area \* 25.

```
max-hops-to-area|mha set <value>
max-hops-to-area|mha [show]
```

**Sets, or displays,** the maximum hops possible in a path to a reachable area. The range for *<value>* is 1 - 30. *<value>* must be greater than or equal to actual max hops to any area.

```
max-hops-to-node|mhn set <value>
max-hops-to-node|mhn [show]
```

**Sets, or displays,** the maximum hops possible in a path to a reachable node. The range for *<value>* is 1 - 30. *<value>* must be greater than or equal to actual max hops in an area.

```
max-node-num|mnn set <value>
max-node-num|mnn [show]
```

**Sets, or displays,** the maximum node number allowed within this area. The range for *<value>* is 1 - 1023. *<value>* must be greater than or equal to maximum node number in use.

```
max-routers|mr pset <value> <seglist>
max-routers|mr [show] [<seglist>]
```

**Sets, or displays,** the number of broadcast router adjacencies supported on the port(s) in *<seglist>*. *<seglist>* is a comma-separated list of ports or **all**. The range for *<value>* is **1 - 20**.

```
hello-time|ht pset <value> <seglist>
hello-time|ht [show] [<seglist>]
```

**Sets, or displays,** the interval for sending hello packets on the port(s) in *<seglist>*. *<seglist>* is a comma-separated list of ports or **all**. The range for *<value>* is **1 - 8191**

```
cost|c pset <value> <seglist>
cost [show] [<seglist>]
```

**Sets, or displays,** the cost for the ports in *<seglist>*. *<seglist>* is a comma-separated list of ports or **all**. The range for *<value>* is **1 - 127**

```
max-visits set <value> <seglist>
max-visits [show]
```

**Sets, or displays,** the maximum visits for a packet before the router assumes that the packet is looping. The range for *<value>* is **maxpath - 60**. *<value>* must be greater than equal to the actual maximum path in the entire network.

```
priority|pri pset <value> <seglist>
priority|pri [show] <seglist>
```

**Sets, or displays,** the priority for the port(s) in *<seglist>*. *<seglist>* is a comma-separated list of ports or **all**. The range for *<value>* is **0 - 127**.

```
update-time|ut set <secs>
update-time|ut [show]
```

**Sets, or displays,** background timer for sending routing updates. The range for *<secs>* is **1 - 1200**.

## 20.2 Segment Configuration

Once the ASN-9000 is configured to forward DECnet packets, you must designate one or more segments as DECnet segments to make the software interpret and forward the correct packets. This step also causes the software to transmit and accept routing control packets over these segments, enabling it to discover neighboring endnodes and routers. There are also several parameters associated with each segment that can be set to tune network performance.

### 20.2.1 Configuration

From the `dec` subsystem prompt, the only necessary segment configuration step is to enable DECnet forwarding for all segments attached to DECnet networks. The `penable dec <seg-list>` command tells the software that DECnet packets may arrive over these segments and that they should be used for routing purposes:

```
penable dec <seglist>
```

This command can be used to either enable or disable DECnet forwarding for each segment. The command uses the normal `<seg-list>` syntax, which is a hyphen- and comma- separated list of segment numbers. For example, if segments 1, 2, and 3 are to be on DECnet networks, the command is:

```
penable dec 2.1-2.4
```

```
193:ASN-9000:dec# penable dec 2.1-2.4
Port 2.1: Okay
Port 2.2: Okay
Port 2.3: Okay
Port 2.4: Okay
```

After enabling the segments, you can verify the segment configuration with the `show priority` command:

```
show priority|pri [<seglist>]
```

For example:

```
196:ASN-9000:dec# dpp 1-2
DECnet port configuration (Port 1)
-----
block-size:      1498
cost:            10
curr-adj-routers: 0
designated-rtr:   aa-00-04-00-1e-8a   (5.1023)
hello-time:      15 seconds
```

## DECnet Subsystem Commands

```
last-hello-sent: 12 seconds ago
mgmt-state:      Enabled
max-routers:    10
priority:       0
run-state:      Up
type:           Ethernet
```

DECnet port configuration (Port 2)

```
-----
block-size:      1498
cost:            10
curr-adj-routers: 0
designated-rtr:   aa-00-04-00-1e-8a  (5.1023)
hello-time:      15 seconds
last-hello-sent: 12 seconds ago
mgmt-state:      Enabled
max-routers:    10
priority:       0
run-state:      Up
type:           Ethernet
```

At this point, you can also verify the routing status of the DECnet software through the **show routing-status** command. This command shows the state of the global DEC forwarding switch as well as whether or not each segment is configured to route DECnet packets.

```
198:ASN-9000:dec# show rs
DECnet routing status:
```

Node/Port	Management State	Routing State
-----	-----	-----
DEC-Forwarding	Enabled	Up
Port 1	Enabled	Up
Port 2	Enabled	Up
Port 3	Enabled	Up
Port 4	Disabled	Down
Port 5	Disabled	Down

*<remaining rows omitted for brevity>*

In this listing, the Management State column refers to DECnet forwarding being enabled or disabled on each segment, while the Routing State column refers to the low-level hardware status. If that segment does not have a cable attached to it (and automatic segment-state detection is enabled), or if the segment has been disabled in the bridging subsystem, the Routing State shows “DOWN” instead of “UP.”

## 20.3 Display Commands

---

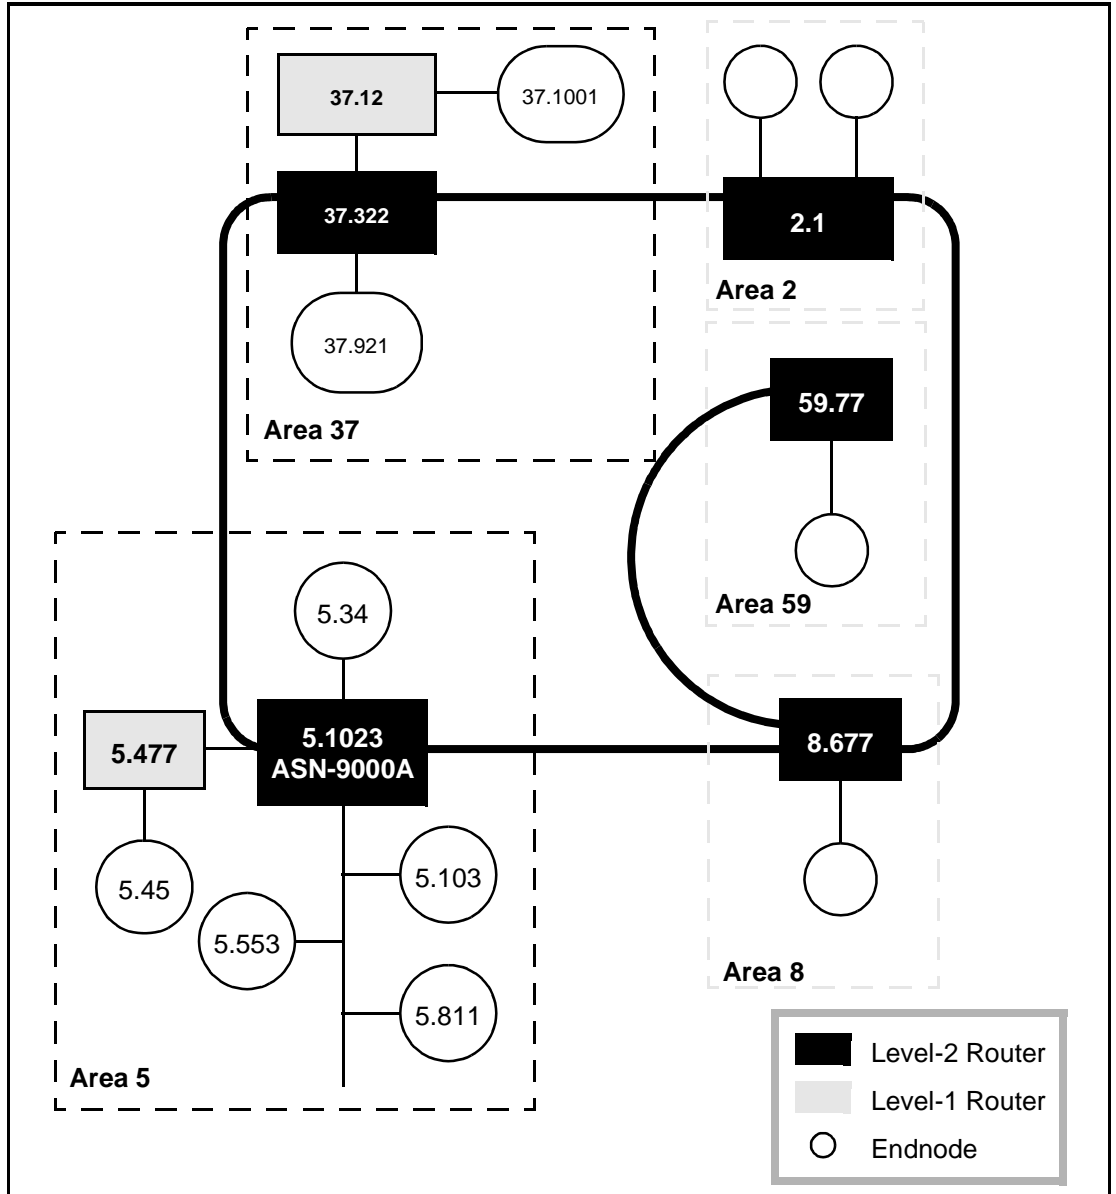
Using the display commands, you can:

- Look for adjacent DECnet routers in the network.
- Look for all DECnet endnodes adjacent to the ASN-9000.
- Look at the DECnet routing table to verify that all the routes are present.

### 20.3.1 Verification of Routing

After the node and segments are configured, the ASN-9000 begins forwarding packets among nearby nodes. To verify that the ASN-9000 has identified its neighbors, you can use one of several display commands to examine routing tables and node lists. Figure 20.2 shows a sample DECnet network. The display commands shown give information about this configuration. Note that we are monitoring the ASN-9000 defined as node 5.1023, located on the left. It is serving as a Level-2 router for area 5, which consists of 7 nodes: itself, 5.34, 5.477, 5.45, 5.103, 5.553, and 5.811. There are 4 other areas, 37, 2, 8, and 59. In Figure 20.2, “endnodes” are depicted as a single circle. (Endnodes, such as non-routing workstations, are nodes not capable of forwarding packets.) Level-1 router nodes are shown as lightly-shaded rectangles.

Note that nodes that are capable of routing, but that appear on the periphery of networks (thus giving them nothing to route to), still qualify as routers and appear on the *ForeRunner* ASN-9000 listings as “routers” rather than “endnodes.” Level-2 router nodes are rectangles, and are connected with bold lines. All connections to the *ForeRunner* ASN-9000 switch are made through the segment numbers listed (1 through 5) by the small digits near the connecting lines. Also note that, while no endnodes are shown on the bold connections (links between departments, for example) between Level-2 routers, the protocols permit them to be there. For example, on the connection between 5.1023 and 37.322, endnodes or Level-1 routers for areas 5 and 37 could be attached. Each Level-2 router would recognize the nodes that belong to its area and forward packets to them.



**Figure 20.2 - Router Verification**

## 20.3.2 Setting and Displaying Block-Size

The block-size command enables you to control the size of internal routing tables. If you have a large network, block-size may need to be raised, but otherwise block-size should be kept low to conserve memory. To set the block-size of the internal routing tables, issue the following command:

```
block-size|bs pset <value> <seglist>
```

To display current block-size, issue the following command:

```
block-size|bs [show] [<seglist>]
```

## 20.3.3 Displaying Adjacent Routers

Look for adjacent routers in the network by typing:

```
adj[acent] [show] r[outer[s]] [[a[ddr]=]<node>]
```

```
407:ASN-9000:dec# show adj r
```

DECnet router adjacency table:

Adj	Node ID	Type	State	Seg	Blk Siz	Hel.Tim	Prior.	Age
1	5.477	Router	Up	2.1	1498	15	0	3
2	37.322	Area Rtr	Up	2.3	1498	15	0	3
3	8.677	Area Rtr	Up	2.5	1498	15	0	3

This command shows all the “adjacent” routers. In DECnet terminology, “adjacent” means “directly connected.” Thus nodes on the other end of 10Base-T links, or other sites on an Ethernet cable, are considered “adjacent.” A “router” is any node which can forward packets. Thus, this command shows all the directly connected routing nodes that the ASN-9000 has discovered. One is inside the ASN-9000 own area (5.477) and the other two are Level-2 routers (“Area Rtr”) in areas 37 and 8.

## 20.3.4 Displaying Adjacent Endnodes

Look for all endnodes adjacent to this router by typing:

```
adj[acent] [show] [end]node[s] [[a[ddr]=]<node>]
```

```
408:ASN-9000:dec# show adj node
```

DECnet end-node adjacency table:

Adj	Node ID	Type	State	Seg	Blk Siz	Hel.Tim	Prior.	Age
1	5.34	End Node	Up	2.1	1498	10	0	9
2	5.811	End Node	Up	2.2	1498	10	0	9
3	5.103	End Node	Up	2.3	1498	10	0	9
4	5.553	End Node	Up	2.4	1498	10	0	9

This command shows all directly connected nodes that are “endnodes,” that is, those which cannot forward packets. The three nodes on the Ethernet cable are endnodes, as is node 5.34 (directly connected to segment 1). Note that node 5.45 is not adjacent, because this *ForeRunner* ASN-9000 switch cannot reach it directly.

## 20.3.5 Displaying the Route Table

Look at the routing tables to verify that all the routes are present by issuing the **display-route-tbl** command. Here is the syntax for this command:

```
show route|rt [<disprestrict>]
```

This command displays the route table, which is maintained by the DECnet routing software. It contains all the routes to nodes in this area that the *ForeRunner* ASN-9000 switch has found dynamically. (DECnet does not provide for static, user-specified routes.)

Here is an example of the display produced by this command:

```
409:ASN-9000:dec# show rt
```

DECnet routing table:

Node	Seg	Next Hop	Hops	Cost
area-rtr	-----	This-Rtr	----	----
5.34	1.1	-----	1	10
5.45	1.2	5.477	2	10
5.103	1.3	-----	1	10
5.477	1.2	-----	1	10
5.553	1.3	-----	1	10
5.811	1.3	-----	1	10
5.1023	Local			

Each entry contains the following information:

<b>Node</b>	The address of the destination node.
<b>Port</b>	The ASN-9000 segment that a packet destined for this node should leave on.
<b>Next Hop</b>	The address of the next node a packet must pass through.
<b>Hops</b>	The number of nodes the packet must pass through.
<b>Cost</b>	A number reflecting the desirability of using this route.

From this table, we see that this area (Area 5) consists of 7 nodes: 34, 45, 103, 477, 553, 811, and 1023. The ASN-9000 is node 1023. The nodes on the Ethernet cable are 103, 553, and 811; they are accessible directly through segment 1.3. Two other nodes, 34 and 477, can be contacted directly through segments 1.1 and 1.2. One node, number 45, can only be reached through node 477, which is a router. Therefore the routing table shows that to send packets to node 45, the “Next Hop” is node 477, and that the node is two hops away from this one.

Note that the ASN-9000, like all DECnet nodes, keeps track of the nearest Level-2 router. Since the ASN-9000 is configured as an area-router (Level-2), the nearest Level-2 router is itself. Consequently, the next hop listed for the “area-rtr” node (the one responsible for all inter-area routing) says “This-Rtr.”

If this router is configured as an area router (Level-2), look at the area table. This is a list of all known areas, along with the best way to get to them. To display this table, issue the following command:

**area [show] [*<area>*]**

```
410:ASN-9000:dec# show area
DECnet area table:
Area  Port    Next Hop                Hops  Cost
----  -
  2    1.4    8.677                    2     20
  5    Local
  8    1.4    8.677                    1     10
 37    1.5    37.322                   1     10
 59    1.4    8.677                    2     20
```

From this example, we can tell that the ASN-9000 is in area 5, and three other areas are accessible through the Level-2 router at 8.677, which is attached to the network on Segment 1.4. The other area is Area 37, available through segment 1.5.

If the ASN-9000 is configured as a Level-1 router (in a multi-area network), the “area-rtr” entry points to another node. As an example, imagine that the other router (node 5.477) is also a ASN-9000.

If we were to examine the route table on that hypothetical node, you would see something like the following:

```
1:OtherASN-9000:dec# show area
```

```
DECnet routing table:
```

Node	Segment	Next Hop	Hops	Cost
-----	----	-----	----	----
area-rtr	1.2	5.1023	1	10
5.34	1.2	5.1023	2	10
5.45	1.1	5.45	1	10
5.103	1.2	5.1023	2	10
5.477	Local			
5.553	1.2	5.1023	2	10
5.811	1.2	5.1023	2	10
5.1023	1.2	5.1023	1	10

Examine the node and segment statistics to verify that the ASN-9000 is receiving data and control packets correctly.

## 20.3.6 Displaying Statistics

There are two types of statistics collected in the DECnet subsystem: node statistics and segment statistics. The node statistics are displayed with the **stats show** command, and they list information that is not associated with any particular segment. All the numbers displayed by the **stats show** command relate to errors or dropped packets, so the ideal display is all zeros. Here is an example of the **stats show** command:

```
411:ASN-9000:dec# stats show n
DECnet node statistics (count since last stats clear):
node unreachable pkt loss      0
aged packet loss               0
node out-of-range pkt loss     0
oversized pkt loss             0
pkt format error               0
partial routing update loss    0
verification reject           0
routing table corrupted        0
no timers for updates          0
no bufs for sending hello      0
invalid hello from router      0
invalid hello from endnode     0
no room for router adj         0
no room for endnode adj        0
low priority rtr bumped        0
no bufs for lvl 1 update       0
lvl 1 msg format error         0
lvl 1 msg checksum error       0
lvl 1 msg area num error       0
no bufs for lvl 2 update       0
lvl 2 msg format error         0
lvl 2 msg checksum error       0
router moved to diff. port     0
end node moved to diff. port   0
```

As in other ASN-9000 subsystems, the DECnet software maintains two copies of the node statistics:

- Count since the last clear.
- Count since the last system reset.

Both counters increment when errors occur, but the **stats clear** command clears only the count since last clear. To display the count since the last reset, use the **-t** option with the **stats show** command.

The segment statistics are collected in the same manner. These statistics are primarily counts of how many DECnet packets are routed through each segment. This can give you an idea of where the most traffic is coming from, and may provide insight on how to better structure the network.

### **20.3.6.1 Displaying the Route Cache**

To display the DECnet route cache, issue the following command:

```
cache [show] [<disprestrict>]
```

To clear the DECnet route cache, issue the following command:

```
cache clear
```

# APPENDIX A

## Configuration Defaults

This appendix lists the configuration defaults for the ASN-9000. The purpose of this appendix is to help understand what is already configured in the software to help aid in diagnosing and troubleshooting the ASN-9000.

### A.1 ASN-9000 Software Subsystems and Defaults

Configuration defaults are listed according to subsystem.

**Table A.1 - Bridge Subsystem**

Command and Description
<b>config [show] [&lt;argument-list&gt; all]</b> Specifies the configuration parameters to display. Default is <b>all</b> .
<b>bt [show] [&lt;seglist&gt; all] [&lt;ethaddr&gt;] [-t] [[-h] [-m]]]</b> Specifies which segment(s) to display bridge table entries. Default is <b>all</b> .
<b>set aging [&lt;time&gt;]off</b> Specifies the aging time to clear learned entries in seconds, complex time (hh:mm:ss) or tiny time (microseconds or milliseconds). Default is 60 minutes.
<b>set aging [&lt;time&gt;]off</b> Specifies the aging time to clear learned entries in seconds, complex time (hh:mm:ss) or tiny time (microseconds or milliseconds). Default is 60 minutes.
<b>enable disable spantree</b> Specifies whether enabling or disabling the Spanning-Tree algorithm. Default is <b>disable</b> .
<b>spantree st set bridge-priorit bp &lt;priority&gt;</b> Specifies the Spanning-Tree bridge priority. Default is 8000 (hex).
<b>spantree st sset seg-priorit sp &lt;priority&gt; &lt;seglist&gt;</b> Specifies the Spanning-Tree segment priority. Default is 8000 (hex).
<b>spantree st sset path-cos pc &lt;path-cost&gt; &lt;seglist&gt;</b> Specifies the cost of the path. Default is 100 for 10Mb/s Ethernet segments, and 10 for Fast Ethernet segments.

**Table A.1 - Bridge Subsystem**

Command and Description
<b>spantree st set maxage &lt;time&gt;</b> Specifies the maximum age, in seconds. Default is 21 seconds
<b>spantree st set hello &lt;time&gt;</b> Specifies the hello time, in seconds. Default is 4 seconds.
<b>spantree st set fwddelay &lt;time&gt;</b> Specifies the forward delay, in seconds. Default is 16 seconds.
<b>spantree st set high-util &lt;percentage&gt;</b> Specifies the upper-end value of segment utilization. This value is a percentage in the range of 1 to 100. Default is 70%.
<b>spantree st set low-util &lt;percentage&gt;</b> Specifies the upper-end value of segment utilization. This value is a percentage in the range of 1 to 100. Default is 50%.

**Table A.2 - Global Commands**

Command and Description
<b>su [root   monitor]</b> Changes the userid to the root or monitor. Default is the root.
<b>rm [-f] [-i] &lt;filespec&gt; [&lt;filespec&gt;...]</b> Overrides the -f (Force) flag, presenting a prompt before removing each file. If not specified, -i is the default.

**Table A.3 - Host Subsystem**

Command and Description
<b>set kadelay   kad &lt;minutes&gt;</b> Specifies how many minutes the ASN-9000 allows a TCP (TELNET) connection to remain idle before sending keep-alive packets. Default is 20 minutes.
<b>set kainterval   kai &lt;seconds&gt;</b> Specifies how often the ASN-9000 sends keep-alive packets before ending a connection. Default is 75 seconds.

Table A.4 - TFTP Subsystem

Command and Description
<b>get -a</b> fore/ph/ethan.env ethan.env Specifies net-ASCII mode. Files are transferred in binary mode by default.
<b>get [-h &lt;host&gt;] [-a] &lt;remote-file&gt; [&lt;local-file&gt;   tty]</b> Specifies the IP address, in dotted-decimal notation, of the TFTP server. If not specified, the default server is used. The default server is specified using the <b>set server</b> command.
<b>put [-h &lt;host&gt;] [-a] &lt;localfile&gt; [&lt;remote-file&gt;]</b> Specifies the IP address, in dotted-decimal notation, of the TFTP server. If not specified, the default server is used. The default server is specified using the <b>set server</b> command.
<b>put [-h &lt;host&gt;] [-a] &lt;localfile&gt; [&lt;remote-file&gt;]</b> Specifies the IP address, in dotted-decimal notation, of the TFTP server. If not specified, the default server is used. The default server is specified using the <b>set server</b> command.

Table A.5 - IP Subsystem

Command and Description
<b>interface add &lt;vlanid&gt; &lt;ipaddr&gt; [&lt;prefixlen&gt;   &lt;mask&gt;] [br[oadcast] 0   1] [met[ric] &lt;metric&gt;]</b> Allows a standard IP subnet mask to be used. If a particular network uses IP subnet addressing, then the subnet mask should be specified here using dotted-decimal notation. Otherwise, the system uses a default subnet mask equal to the “natural” subnet mask for the particular class of address. <b>[br[oadcast] 0   1]</b> Specifies the style of broadcast address on a segment-by-segment basis. The default is <b>br1</b> . <b>[met[ric] &lt;metric&gt;]</b> Specifies an additional cost of using the subnet interface. The default is zero.
<b>route enable   disable &lt;destination&gt; &lt;gw-ipaddr&gt; &lt;metric&gt; &lt;segment&gt;</b> Specifies whether you are enabling or disabling IP routing. The default is disable.
<b>arp set   show   unset age &lt;time&gt;</b> <time> Specifies (in minutes) a new aging interval or turns aging off. The default is 5 minutes.

Table A.5 - IP Subsystem

Command and Description
<p><b>ping   pi [-t &lt;timeout&gt;] [-size &lt;size&gt;] &lt;ipaddr&gt;</b></p> <p><b>[-t &lt;timeout&gt;]</b> Specifies how many seconds the <i>ForeRunner</i> ASN-9000 switch waits for a response from the specified device. The default is 5 seconds.</p> <p><b>[-size &lt;size&gt;]</b> Specifies the packet length. You can specify any length from 64 through 1472 bytes. The default is 64 bytes.</p>
<p><b>ipdefaultttl   ittl set &lt;value&gt;</b> Specifies the new TTL time in hops. The default is 16 hops.</p>
<p><b>enable   disable send-icmp-redirect   sir</b> Specifies whether you are enabling or disabling ICMP redirect messages. The default is <b>enable</b>.</p>
<p><b>enable   disable fwd-pkts-with-srcrt-option   fps</b> Specifies whether you are enabling or disabling source-route filtering. The default is <b>enable</b>.</p>

Table A.6 - IP Multicast Subsystem

Command and Description
<p><b>it   interface add &lt;ipaddr&gt; [met[ric]&lt;metric&gt;] [thresh[old]&lt;thresh&gt;]</b> <b>[met[ric]&lt;metric&gt;]</b> Specifies an additional cost (measured in hops to the destination) of using the interface. The default is 1.</p> <p><b>[thresh[old]&lt;thresh&gt;]</b> Specifies the minimum time-to-live (TTL) value that an IP Multicast packet must have before it is forwarded over this interface. The default is 1.</p>

Table A.6 - IP Multicast Subsystem

Command and Description
<b>tunnel add [-s] loc[al]&lt;local-addr&gt; rem[ote]&lt;remote-addr&gt; [met[ric]&lt;mv&gt; [thresh[old]&lt;tv&gt;]</b> <b>[met[ric]&lt;mv&gt;</b> Specifies an additional cost (extra hops to the destination) of using the virtual interface with which this tunnel is associated. The default is 1. <b>[thresh[old]&lt;tv&gt;]</b> Specifies the minimum time-to-live (TTL) value that an IP Multicast packet must have before it can be forwarded through the tunnel. The default is 1.
enable   disable ipm Specifies whether enabling or disabling IP Multicast forwarding. The default is <b>disable</b> .
<b>penable   pdisable transmit &lt;segment-list&gt;</b> Specifies whether enabling or disabling IP Multicast forwarding. The default is <b>penable</b> .
enable   disable multicast-aware-bridging Specifies whether enabling or disabling multicast-aware-bridging. The default is <b>disabled</b> .
<b>it   interface add &lt;ipaddr&gt; [met[ric]&lt;metric&gt;] [thresh[old]&lt;thresh&gt;]</b> <b>[met[ric]&lt;mv&gt;</b> Specifies an additional cost (measured in hops to the destination) of using the interface. The default is 1. <b>[thresh[old]&lt;tv&gt;]</b> Specifies the minimum time-to-live (TTL) value that an IP Multicast packet must have before it can be forwarded over this interface. The default is 1.
enable   disable fwd-pkts-with-srcrt-option   fps Specifies whether enabling or disabling source-route filtering. The default is <b>enable</b> .

Table A.7 - IP/RIP Subsystem

Command and Description
rip-bridging   rb [enable   disable] Enables or disables the RIP bridging feature. The default is <b>disable</b> .

**Table A.8 - IP/OSPF Subsystem**

Command and Description
<p>asbd enable   disable</p> <p>Specifies whether to enable or disable Autonomous System Border router. The default is <b>disable</b>.</p>
<p>auto-vlink enable   disable</p> <p>Specifies whether to enable or disable the automatic virtual-link feature. The default is <b>enable</b>.</p>
<p>area add &lt;area-id&gt; [&lt;auth-type&gt;] [stub-area-cost   sac &lt;cost&gt;]</p> <p>&lt;auth-type&gt; <b>md5</b>   <b>m</b></p> <p>Specifies that MD5 authentication is required for OSPF packets sent within this area. The default is none (no authentication).</p> <p>stub-area-cost   sac &lt;cost&gt;</p> <p>The OSPF software configures the default route automatically.</p>

**Table A.9 - ATALK Subsystem**

Command and Description
<p>enable   disable atalk</p> <p>Specifies whether to enable or disable AppleTalk routing. The default is <b>disable</b>.</p>
<p>ping [-t &lt;timeout&gt;] [-size &lt;pktsize&gt;] &lt;net&gt;.&lt;node&gt;</p> <p>[-t &lt;timeout&gt;]</p> <p>Optionally specifies the number of seconds the ASN-9000 waits to receive a reply packet from the specified node. The default is <b>15</b> seconds.</p> <p>[-size &lt;pktsize&gt;]</p> <p>If the &lt;timeout&gt; argument is used, optionally specifies the size of the echo packet to send to the node. The default is <b>64</b> bytes.</p>

**Table A.10 - IPX Subsystem**

Command and Description
<b>interface   it add</b> <segmentlist> <network> <b>[mtu</b> <mtu>] <b>[met[ric]</b> <metric>] <b>[encap enet   802.3   802.2   snap]</b> <b>[mtu</b> <mtu>] Specifies the maximum transmission unit (number of octets) for packets forwarded on this segment.
<b>enable   disable [ipx]</b> Specifies whether enabling or disabling IPX forwarding. The default is disable.
<b>set ripsap-ctrl   rsct [normal   n vlan   v]</b> <b>normal   n</b> Specifies that RIP and SAP updates are generated on a per-segment basis. This is the default.
<b>penable   pdisable type20-port-forwarding   tpfw</b> <seglist> Specifies whether to enable or disable type-20 packet forwarding. The default is penable (enabled).

**Table A.11 - DECnet Subsystem**

Command and Description
<b>set max-node-num   mnn</b> <value> This determines the number of nodes that can exist within the ASN-9000 . The default is 255.



# APPENDIX B

## Well-Known Ports

This appendix lists the well-known names provided in RFC 1340 that the *ForeRunner* ASN-9000 system supports. When configuring an ASN-9000 IP or TCP filter, either the port number or the well-known name can be supplied to specify the destination port of packets to either block or accept. Supply the port number or well-known name in the *<dstseg>* field of templates for any TCP and IP filters created. The *<dstseg>* field is used with the following TCP and IP filter commands:

- tcp tcp-filter add
- ip ip-fil-acsc-ctrl add

When an IP packet comes in on a segment, the packet then on the IP header to begin routing it through the network to its eventual destination. In the IP header, the protocol type field denotes the kind of packet that follows, such as ARP, TCP, or UDP.

If the protocol type field indicates a TCP or UDP packet, a 16-bit number represents the source and destination ports. Many of these ports are considered “well-known” ports because they appear in an official, published table (RFC 1340) that relates the names of commonly-used protocols with the TCP or UDP ports they typically use.

Table B.1 lists the well-known port names that the ASN-9000 recognizes, and provides the port number associated with each well-known name. Enter the “well-known” port name or number exactly as shown in the table.

**Table B.1 - Well Known Names and Ports**

Well-known Name	Port Number	Well-known Name	Port Number
at-echo	204	at-nbp	202
at-rtmp	201	at-zis	206
auth	113	bgp	179
biff	512	bootpc	68
bootps	67	chargen	19
courier	530	csnet-ns	105
daytime	13	discard	9
dls	197	domain	53
echo	7	exec	512
finger	79	ftp	21
ftp-data	20	hostname	101
hostnames	101	ingreslock	1524
ipcserv	600	ipx	213
iso-tp0	146	isop-tsap	102
kerberos	88	klongin	543
kshell	544	link	87
login	513	lpd	515
monitor	561	nameserver	42
netbios-dgm	138	netbios-ns	137
netbios-ssn	139	netnews	532
netstat	15	news	144
NeWs	144	new-rwho	550
nicname	43	nnntp	119
npp	92	ntp	123
pcserver	600	pop-2	109

**Table B.1** - Well Known Names and Ports

Well-known Name	Port Number	Well-known Name	Port Number
pop3	110	printer	515
print-srv	170	rip	520
rlogin	513	rmonitor	560
route	520	rtelnet	107
rwho	513	shell	514
smtp	25	snmp	161
snmptrap	162	sunrpc	111
supdup	95	syslog	514
systat	11	tacnews	98
talk	517	tcpmux	1
telnet	23	tftp	69
time	37	timed	525
uucp	540	who	513
whois	43	x400	103
x400-snd	104	xdmcp	177
supdup	95	syslog	514

## *Well-Known Ports*

# Index

## Symbols

!	5 - 1, 9 - 3
!!	5 - 1, 9 - 3
^	5 - 1

## A

AARP table	
description	17 - 15
adding	
static ARP entry	13 - 24
address	
host	13 - 4
IP	
restrictions	13 - 5
IP interface	
displaying	13 - 6
adjacent router	20 - 13
ADVRT-PRINT	18 - 20
aging	
bridge table	
description	11 - 8
algorithm	
IP route cache	13 - 40, 14 - 12
alias	
environment parameter	5 - 9
loading	5 - 5
all-0s broadcast address	13 - 7, A - 3
all-1s broadcast address	13 - 7
allocating memory	
DECnet	20 - 2
AppleTalk	
Routing Table Maintenance	
Protocol (RTMP)	17 - 17

AppleTalk Address Resolution	
Protocol (AARP)	17 - 15
AppleTalk route table	
displaying	17 - 17
AppleTalk subsystem	17 - 1
ARCHIVE-SVR	18 - 20
area	
adding	16 - 9
adding a network range	16 - 13
deleting	16 - 11
deleting a network range	16 - 14
area del command	16 - 11
area number	20 - 2
ARP	
proxy	13 - 36
ARP requests	13 - 21
ARP table	
aging	
setting	13 - 23, 17 - 17
description	13 - 21
arp-tableclear command	17 - 17
automatic segment-state detection	8 - 6

## B

Backspace key	3 - 3
baud rates	
displaying	6 - 6
setting	6 - 2, 6 - 16
boot messages	
example	4 - 2
boot source	
displaying	6 - 18
booting	6 - 12

## Index

- troubleshooting ..... 4 - 3, 6 - 9
- bridge
  - configuration
    - displaying ..... 11 - 1
- bridge cache
  - clearing ..... 11 - 19
  - description ..... 1 - 5, 11 - 18
  - displaying ..... 11 - 19
- bridge state command ..... 8 - 7
- bridge table
  - clearing ..... 11 - 6
  - displaying ..... 11 - 4
  - flags ..... 11 - 6
- bridging ..... 19 - 1
- broadcast ID ..... 13 - 4
- build\_oid ..... 12 - 6
- C**
- canceling
  - line of input ..... 3 - 3
- card-swap command ..... 6 - 3
- character
  - wildcard ..... 13 - 23
- chassis
  - displaying or setting the system
    - location description ..... 6 - 14
  - system name
    - displaying or setting ..... 6 - 14
- clearing
  - ARP table ..... 13 - 23
  - statistics
    - AppleTalk ..... 17 - 23
- command prompt
  - runtime ..... 3 - 1
- commands
  - area del ..... 16 - 11
  - arp-table (ip) ..... 13 - 22
  - arp-tableclear (ip) ..... 13 - 23
  - bye ..... 4 - 5
  - card-swap ..... 6 - 3
  - date ..... 6 - 6
  - dec rs ..... 20 - 2
  - del-interface (ip) ..... 13 - 9
  - del-route net ..... 13 - 18
  - display-area-tbl ..... 20 - 15
  - display-cache (bridge) ..... 11 - 19
  - display-port-param ..... 20 - 9
  - display-routecache (ip) ..... 13 - 39
  - display-router-adj ..... 20 - 13
  - flush-cache (bridge) ..... 11 - 19
  - flush-routecache (ip) ..... 13 - 40, 14 - 12
  - fremove ..... 5 - 3
  - getmem ..... 14 - 2, 20 - 2
  - idprom ..... 6 - 8
  - interface-table (ip) ..... 13 - 6
  - listdir ..... 5 - 3
  - logout ..... 4 - 5
  - main getmem ..... 16 - 2
  - nvramp set ..... 7 - 2
  - passwd ..... 6 - 9
  - password ..... 6 - 9
  - ping (ip) ..... 13 - 25
  - proxy-arp (ip) ..... 13 - 37
  - readcfg ..... 6 - 11
  - reboot ..... 6 - 12
  - remove ..... 5 - 3
  - route-table ..... 14 - 10
  - route-table (ip) ..... 13 - 13
  - route-table (ipx) ..... 18 - 8
  - savecfg ..... 6 - 12
  - saveenv ..... 5 - 11
  - set asbd ..... 16 - 4, A - 6
  - set bridge-net-bcast ..... 13 - 36
  - set bridge-priority ..... 11 - 14
  - set ipDefaultTTL ..... 13 - 33

- set ipxFowarding ..... 18 - 6, A - 7
- set route-net-bcast ..... 13 - 36
- set router-id ..... 16 - 3, 16 - 4
- set seg-priority ..... 11 - 15
- set timer-threshold ..... 11 - 16
- set-arpge (ip) ..... 13 - 23
- set-node-param node-type
  - area-router ..... 20 - 2
- show lsdb ..... 16 - 17
- show neighbor ..... 16 - 15
- show net-range ..... 16 - 14
- show virtual-link ..... 16 - 22
- showcfg (ip) ..... 13 - 2
- showcfg (ipm) ..... 14 - 3
- show-mcast-groups ..... 14 - 4
- show-neighbors ..... 14 - 4
- spantree ..... 11 - 14
- stats ..... 16 - 26, 16 - 27
- stats (bridge) ..... 11 - 18
- stats (ip) ..... 13 - 38
- stats (rip) ..... 15 - 7, 15 - 8
- stats-clear (bridge) ..... 11 - 18
- stats-clear (ip) ..... 13 - 39
- stats-clear (rip) ..... 15 - 7
- stats-clear dvm | igmp | rt
  - 14 - 15
- syslocn ..... 6 - 14
- sysname ..... 6 - 14
- temperature ..... 6 - 15
- timedcmd add ..... 5 - 7
- timedcmd off ..... 5 - 2, 5 - 8
- timedcmd on ..... 5 - 8
- virtual-link add ..... 16 - 5, 16 - 21
- virtual-link del ..... 16 - 20, 16 - 21
- configuration
  - default
    - ARP ..... 13 - 23
  - IP
    - displaying ..... 13 - 2
  - RIP
    - displaying ..... 18 - 17
    - per-VLAN ..... 15 - 1
  - SAP
    - displaying ..... 18 - 17
    - saving ..... 6 - 12
  - configuration file
    - editing ..... 6 - 19
    - loading from a TFTP server .10 - 8, 10 - 9
    - reading ..... 6 - 10
    - saving ..... 6 - 12
    - saving to a TFTP server ..... 10 - 9
  - configuration requirements
    - IPX translation bridging ..... 19 - 2
  - configuring
    - hub
      - DECnet node ..... 20 - 5
    - IPX RIP ..... 18 - 13
    - IPX SAP ..... 18 - 13
  - connection idle time
    - discussion ..... 9 - 5
  - cost ..... 13 - 8
  - stub area
    - displaying ..... 16 - 12
  - counters
    - statistics ..... 13 - 38
  - crash-reboot
    - setting ..... 7 - 1
  - creating a virtual LAN ..... 13 - 4
  - CTRL + H ..... 3 - 3
  - CTRL + Q ..... 3 - 3
  - CTRL + S ..... 3 - 3
- D
  - date
    - displaying or setting ..... 6 - 6

date command .....	6 - 6	route table (DECnet) .....	20 - 14
DEC routing		statistics .....	20 - 17
setting up .....	20 - 1	IPX, type-20 .....	18 - 24
DECnet		TCP table .....	9 - 3
network		UDP table .....	9 - 7
topology restrictions .....	20 - 4	display-router-adj command .....	20 - 13
DECnet routing .....	20 - 1	Domain Name System .....	13 - 26
default		DRAM .....	6 - 6
network group .....	11 - 9	dynamic entries	
defining		bridge table	
network group .....	11 - 9	description .....	11 - 4
Delete key .....	3 - 3	<b>E</b>	
deleting		enabling	
address		IP Multicast traffic .....	14 - 10
IPX Helper .....	18 - 23	IPX Helper .....	18 - 22
network group .....	11 - 11	enabling or disabling	
static ARP entry .....	13 - 25	IP forwarding .....	13 - 12, 13 - 17
detection		encapsulation	
segment-state .....	8 - 6	IPX bridging .....	19 - 1
deviation number		Enter key .....	3 - 3
reading .....	6 - 9, 6 - 17	environment file	
directory		example .....	5 - 10
displaying .....	5 - 3	EraseChar .....	3 - 3
dispcfg .....	6 - 19	erasing	
display commands .....	20 - 11	characters .....	3 - 3
display-endnode-adj command .....	20 - 14	<b>F</b>	
displaying		file	
AppleTalk Address Resolution		configuration	
Protocol (AARP) table ....	17 - 15	editing .....	6 - 19
AppleTalk configuration .....	17 - 3	FILE-SERVER .....	18 - 20
ARP table .....	13 - 22	flag	
DECnet routing status .....	20 - 10	broadcast .....	13 - 22
DECnet statistics .....	20 - 17	system .....	13 - 22
IP configuration .....	13 - 2	Flash Memory Module	
IPX Helper .....	18 - 23	displaying volume information ...	5 - 3
IPX RIP configuration .....	18 - 17	floppy diskette	
route table .....	18 - 8		

displaying volume information ... 5 - 3  
 ForeView Network Management ..... 1 - 8  
 forwarding  
     DECnet  
         enabling ..... 20 - 9  
**G**  
 gateway  
     IPX ..... 18 - 9  
 groups, IP Multicast  
     displaying ..... 14 - 4  
**H**  
 history  
     displaying ..... 5 - 1  
 host  
     static route  
         defining ..... 13 - 15  
**I**  
 ICMP ..... 13 - 25  
 ICMP echo request packet ..... 13 - 25  
 ID PROM  
     reading ..... 6 - 8  
 idprom command ..... 6 - 8  
 IEN-116 Name Server ..... 13 - 26  
 interface  
     AppleTalk  
         adding ..... 17 - 8  
         deleting ..... 17 - 13  
     IP  
         adding ..... 13 - 6  
         deleting ..... 13 - 9  
     IPX  
         deleting ..... 18 - 5  
 interface table command ..... 17 - 10  
 interfaces  
     IPX  
         adding ..... 18 - 3

interface-table ..... 14 - 6  
 IOP ..... 6 - 6  
 IP  
     broadcast ID ..... 13 - 4  
     host number ..... 13 - 4  
 IP address ..... 18 - 9  
 IP address (gateway)  
     setting ..... 7 - 1  
 IP address (PowerHub 7000 as boot client)  
     setting ..... 7 - 1  
 IP address (TFTP file server)  
     setting ..... 7 - 1  
 IP Helper  
     adding an address ..... 13 - 28, 13 - 31  
     deleting an address ..... 13 - 31, 13 - 32  
     discussion ..... 13 - 26, 13 - 27  
     statistics  
         clearing ..... 13 - 30  
         displaying ..... 13 - 29  
 IP Multicast  
     deleting a tunnel ..... 14 - 7, 14 - 9  
     displaying the interface table ..... 14 - 6  
     displaying the route table ..... 14 - 10  
     enabling ..... 14 - 9  
     enabling traffic ..... 14 - 10  
     flushing the route table ..... 14 - 12  
 IP network number ..... 13 - 4  
 IP parent network ..... 13 - 5  
 IP route  
     deleting ..... 13 - 18  
 IP route cache  
     description ..... 13 - 39, 14 - 12  
 IP routing  
     statistics ..... 13 - 6, 14 - 4  
 IPX  
     server table ..... 18 - 18  
 IPX control type

## Index

- per-segment ..... 18 - 13
  - per-VLAN ..... 18 - 13
- IPX Helper
  - enabling ..... 18 - 22
- IPX RIP
  - discussion ..... 18 - 17
- IPX routing
  - relationship to IPX translation bridging ..... 19 - 1
- IPX SAP
  - discussion ..... 18 - 17
- IPX translation bridging ..... 19 - 1
  - configuration requirements ..... 19 - 2
  - enabling or disabling ..... 19 - 2
  - relationship to IPX routing ..... 19 - 1
- IPX translation-bridging network
  - adding ..... 19 - 2
  - deleting ..... 19 - 4
  - displaying ..... 19 - 3
- issuing commands ..... 3 - 1
- J**
- JOB-SERVER ..... 18 - 20
- K**
- keep-alive interval
  - discussion ..... 9 - 5
- Kermit
  - loading a command alias ..... 5 - 6
- L**
- learned addresses
  - AppleTalk Address Resolution Protocol (AARP) table .... 17 - 15
  - ARP table ..... 13 - 21
- Level-1 router ..... 20 - 4
- Level-2 router ..... 20 - 2, 20 - 4
- link
  - virtual
    - adding ..... 16 - 5
    - deleting ..... 16 - 21
  - logging in ..... 4 - 4, 4 - 6
  - logging out ..... 4 - 5
  - login
    - setting a password ..... 6 - 9
  - login prompt ..... 4 - 4, 4 - 6
- M**
- main getmem command ..... 16 - 2
- main memory ..... 6 - 6
- management capability
  - command prompt ..... 3 - 2
  - password ..... 6 - 9
  - setting a password ..... 6 - 9
- managers
  - SNMP
    - adding ..... 12 - 4
- manual entries
  - bridge table
    - description ..... 11 - 4
- maximum transmission unit (mtu) 18 - 4, A - 7
- max-node-num parameter ..... 20 - 5
- memory allocation
  - AppleTalk ..... 17 - 2
  - DECnet ..... 14 - 2, 20 - 2
  - IPX ..... 18 - 1
- mib2schema ..... 12 - 6
- mixed media
  - IPX translation bridging ..... 19 - 1
- model number
  - reading ..... 6 - 8, 6 - 17
- monitor.env ..... 5 - 11
- N**
- Name Binding Protocol (NBP) ..... 17 - 21
- natural subnet mask ..... 13 - 7, A - 3
- neighbor

- displaying information ..... 16 - 15
- neighbors, IP Multicast
  - displaying ..... 14 - 4
- net
  - directly-attached ..... 13 - 15
- NetBIOS Datagram Server ..... 13 - 26
- NetBIOS Name Server ..... 13 - 26
- network
  - hiding ..... 16 - 13
- network address
  - testing ..... 17 - 23
- network group
  - discussion ..... 11 - 9
- network number ..... 13 - 4
- network range ..... 16 - 13
  - deleting ..... 16 - 14
  - displaying information ..... 16 - 14
- network topology
  - restrictions
    - DECnet ..... 20 - 4
- NIM (Network Interface Module)
  - ID PROM
    - reading ..... 6 - 8
  - live insertion ..... 6 - 3
  - swapping ..... 6 - 3
  - temperature
    - reading ..... 6 - 15
- node ID
  - DECnet ..... 20 - 2
- non-seed segment ..... 17 - 8
- NVRAM
  - variable
    - setting ..... 7 - 1
- nvram set command ..... 7 - 2
- O**
  - object identify (OID) file ..... 12 - 6
  - old\_default ..... 11 - 10
- OSPF area
  - adding ..... 16 - 9
  - adding a network range ..... 16 - 13
  - deleting ..... 16 - 11
  - deleting a network range ..... 16 - 14
- OSPF router ID
  - assigning ..... 16 - 2
- ospf subsystem
  - accessing ..... 16 - 1
- P**
  - Packet Accelerator ..... 6 - 6
  - Packet Engine
    - ID PROM
      - reading ..... 6 - 8
    - reset switch ..... 4 - 1
    - temperature
      - reading ..... 6 - 15
  - parameter
    - max-node-num ..... 20 - 4, 20 - 5
    - node type ..... 20 - 5
  - parent network ..... 13 - 5
  - passwd command ..... 6 - 9
  - password
    - adding ..... 16 - 9
    - changing ..... 6 - 9
    - forgotten ..... 6 - 10
    - setting ..... 6 - 9
  - permanent addresses
    - AppleTalk Address Resolution Protocol (AARP) table .... 17 - 15
    - ARP table ..... 13 - 21
  - physical interface, IP Multicast
    - deleting ..... 14 - 7
  - ping
    - ip subsystem ..... 13 - 25
    - Unix ..... 13 - 25
  - power requirements

- reading from ID PROM ..... 6 - 9
- PRINT-QUEUE ..... 18 - 20
- PRINT-SERVER ..... 18 - 20
- PROM
  - ID
    - reading ..... 6 - 8
- public community ..... 12 - 3
- R**
  - readcfg command ..... 6 - 11
  - reboot command ..... 6 - 12
  - rebooting ..... 6 - 12
  - REM-BRIDGE ..... 18 - 20
  - reset switch ..... 4 - 1
  - restrictions
    - network topology
      - DECnet ..... 20 - 4
  - RFC 1027 ..... 13 - 36
  - RIP ..... 15 - 1
  - RIP (Routing Information Protocol) .... 15 - 1
  - ripsap-ctrl-type command 18 - 13, 18 - 17, A - 7
  - root capability
    - setting a password ..... 6 - 9
  - root.env ..... 5 - 11
  - route cache
    - displaying
      - IPX route cache ... 17 - 20, 18 - 10
    - IP
      - displaying ..... 13 - 39, 14 - 12
      - flushing ..... 13 - 40, 14 - 12
    - IPX
      - displaying ..... 17 - 20, 18 - 10
      - flushing ..... 17 - 20, 18 - 12
  - route cache, flushing ..... 17 - 20, 18 - 12
  - route state ..... 18 - 21
    - IPX ..... 18 - 9
  - route table ..... 13 - 8
    - DECnet
      - displaying ..... 20 - 14
- router
  - IP ..... 13 - 1
- router ID
  - assigning ..... 16 - 2
- routing
  - IP
    - statistics ..... 13 - 6
  - Level-2 ..... 20 - 5
  - verifying ..... 20 - 11
- RS-232 (TTY) port
  - reading a configuration ..... 6 - 10
  - saving a configuration ..... 6 - 12
  - setting baud rate ..... 6 - 2, 6 - 16
- S**
  - savecfg command ..... 6 - 12
  - saving AppleTalk configuration ..... 17 - 4
  - schema file ..... 12 - 6
  - scroll
    - environment parameter ..... 5 - 9
  - segment
    - automatic state detection ..... 8 - 6
    - DECnet
      - configuring ..... 20 - 9
  - segment state
    - detection methods ..... 8 - 6
  - segment-state detection
    - setting ..... 8 - 6
  - serial number
    - reading ..... 6 - 8, 6 - 17
  - server ..... 18 - 21
    - TFTP
      - displaying default ..... 10 - 5
      - downloading or displaying
        - a file ..... 10 - 6
      - setting default ..... 10 - 5

- unsetting default ..... 10 - 5
- server name ..... 18 - 21, 18 - 22
- server table ..... 18 - 18
- server, static ..... 18 - 18
- set asbd command ..... 16 - 4, A - 6
- set router-id command ..... 16 - 3, 16 - 4
- set type20-forwarding command ..... 18 - 25
- setbaud command ..... 6 - 2, 6 - 16
- set-node-param node-type area-router
  - command ..... 20 - 2
- set-port-param mgmt-state enl command 20 - 6
- setting
  - bridge priority ..... 11 - 14
  - segment priority ..... 11 - 15
  - timer threshold ..... 11 - 16
- show lsdb command ..... 16 - 17
- show neighbor command ..... 16 - 15
- show net-range command ..... 16 - 14
- SNMP
  - default configuration ..... 12 - 3
  - managers
    - deleting ..... 12 - 5
- space
  - displaying for Flash Memory Module
    - or floppy diskette ..... 5 - 3
- Spanning-Tree algorithm
  - enabling or disabling ..... 11 - 13
- static entries
  - bridge table
    - adding ..... 11 - 6, 11 - 8
- statistics
  - AppleTalk Echo Protocol (AEP)
    - displaying ..... 17 - 21
  - bridge
    - clearing ..... 11 - 18
  - displaying
    - type-20 and IPX ..... 18 - 24

- ICMP
  - clearing ..... 13 - 39
  - displaying ..... 13 - 38
- IP Multicast
  - clearing ..... 14 - 15
  - displaying ..... 14 - 12
- IPX
  - clearing ..... 18 - 24
- RIP
  - displaying ..... 15 - 7
- SNMP
  - displaying ..... 12 - 2
- TCP
  - clearing ..... 9 - 5
- stats (ipm) command ..... 14 - 12
- stats command ..... 16 - 26, 16 - 27
- stats-clear (ipm) ..... 14 - 15
- stats-clear command ..... 17 - 23
- stub area
  - configuring ..... 16 - 9
- subnet addressing ..... 13 - 7, A - 3
- subnet mask ..... 13 - 7, A - 3
- subnet mask (PowerHub 7000 as boot client)
  - setting ..... 7 - 1
- subsystem
  - command prompt ..... 3 - 2
  - RIP ..... 15 - 1
  - SNMP ..... 12 - 1
  - TFTP ..... 10 - 1
- swapping a NIM ..... 6 - 3
- switch
  - reset ..... 4 - 1
- syslocn command ..... 6 - 14
- sysname command ..... 6 - 14
- system name
  - command prompt ..... 3 - 2
  - displaying or setting ..... 6 - 14

## T

TACACS service ..... 13 - 26

TCP ..... 9 - 1

- configuration
- displaying ..... 9 - 2

Technical Support

- contacting ..... iii

TELNET

- statistics ..... 9 - 4

temperature

- sensor
- reading ..... 6 - 15

temperature command ..... 6 - 15

testing

- network address ..... 17 - 23

TFTP

- UDP port ..... 13 - 26

time

- displaying or setting ..... 6 - 6

Time service ..... 13 - 26

timed command

- environment parameter ..... 5 - 9

timed commands ..... 5 - 6

time-to-live (TTL) parameter

- Setting ..... 13 - 33

Transmission Control Protocol (TCP) .... 9 - 1

trap ..... 12 - 6

traps

- SNMP ..... 12 - 4
- displaying ..... 12 - 2

TTY (RS-232) port

- baud rate
- setting ..... 6 - 2, 6 - 16
- reading a configuration ..... 6 - 10
- saving a configuration ..... 6 - 12

## U

UDP ..... 12 - 1, 15 - 1

## V

variable-length subnet addresses ..... 13 - 5

verifying

- routing ..... 20 - 11

virtual link

- adding ..... 16 - 5, 16 - 20
- deleting ..... 16 - 21
- displaying information ..... 16 - 22

virtual-link add command ..... 16 - 5, 16 - 21

virtual-link del command ..... 16 - 20, 16 - 21

virtual-link feature

- enabling ..... 16 - 5

## W

workstation

- remote ..... 13 - 4

## Z

zone

- adding ..... 17 - 4
- displaying ..... 17 - 6